Virtual Machine/ Enterprise Systems Architecture Version 2 Release 1.0 Performance Report

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 2.1.0 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 2.1.0. See the IBM Programming Announcement for VM/ESA 2.1.0 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:

AIX
ACF/VTAM
CICS
CICS/VSE
ECKD
Enterprise System/9000
ES/9000
IBM
OfficeVision
OpenEdition
OS/2
PR/SM
Processor Resource/Systems Manager
RACF
RAMAC
RS/6000
SAA
Streamer
System/390
S/390
Virtual Machine/Enterprise Systems 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

Acknowledgements

The following people contributed to this report:

Glendale Programming Laboratory

Bill Bitner
Wes Ernsberger
Greg Gasper
Bill Guzior
Larry Hartley
Gary Hine
Raine June
Greg Kudamik
Tom Wright

Washington Systems Center

Marty Horan
Chuck Morse

PC Server S/390 Competency Center in Atlanta

Gary Eheman

Abstract

The *VM/ESA Version 2 Release 1.0 Performance Report* summarizes the performance evaluation of VM/ESA 2.1.0. 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 2.1.0 performance evaluations conducted by the Glendale Programming Laboratory.

Discussion concentrates on the performance changes in VM/ESA 2.1.0, the performance effects of migrating from VM/ESA 1.2.2 to VM/ESA 2.1.0, and the performance of new functions provided in VM/ESA 2.1.0. A number of additional evaluations are also included.

Referenced Publications

The following publications are referred to in this report.

- *VM/ESA: Performance,* SC24-5782
- *VM/ESA: CMS File Pool Planning, Administration, and Operation,* SC24-5751
- *VM/ESA: CP Diagnosis Reference,* LY24-5256
- *PC Server 500 System/390 Performance*, WSC Flash 9522, G023542 or PCSVR390 PACKAGE on MKTTOOLS
- *IBM PC Server 500 S/390 ...Is it right for you?*, GK20-2763 or PCSVR390 PACKAGE on MKTTOOLS
- *VSE/ESA 2.1.0 Performance Considerations*, VE21PERF PACKAGE on IBMVSE
- *VSE/ESA 2.1 Turbo Dispatcher Performance*, VE21PERF PACKAGE on IBMVSE
- *IBM RAMAC Array Family*, GG24-2509
- *Using the IBM RAMAC Array DASD in an MVS, VM, or VSE Environment*, GC26-7013
- *TCP/IP Version 2 Release 2 for VM: Planning and Customization*, SC31-6082

The following publications are performance reports for earlier VM/ESA releases.

- *VM/ESA Release 1.0 Performance Report,* ZZ05-0469[1]
- *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

[1] This report 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.

Summary of Key Findings

This report summarizes the performance evaluation of Virtual Machine/Enterprise Systems Architecture* (VM/ESA*) Version 2 Release 1.0. 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 2.1.0 includes a number of performance enhancements (page 9). Some changes have the potential to adversely affect performance, especially in storage constrained CMS environments (page 16). Lastly, a number of changes were made that benefit VM/ESA performance management (page 17).

Migration from VM/ESA 1.2.2: Benchmark measurements show the following performance results for VM/ESA 2.1.0 relative to VM/ESA 1.2.2:

CMS-intensive Performance has been improved significantly. Benchmark results show internal throughput rate (ITR) improvements of 3.8% to 7.0% and external response time improvements of 7% to 20% (page 25).

The use of compiled REXX by CMS is the key factor that resulted in these improvements. Measurement results indicate that systems that already compile the S-disk REXX execs and XEDIT macros may experience a slight decrease in performance when migrating from VM/ESA 1.2.2 to VM/ESA 2.1.0 (page 31).

VSE guest ITR and elapsed times are equivalent for the DYNAPACE I/O-intensive batch workload (page 47). These comparisons include results on the 9121-480 2-way processor using the VSE/ESA* 2.1.0 Turbo Dispatcher. ITR and response times are equivalent for the VSECICS transaction processing workload (page 64).

VMSES/E The performance of the VMFBLD function has been improved. Elapsed time reductions ranging from 10% to 24% were observed (page 70).

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 1.1.5 (370 Feature) or VM/ESA 2.1.0 (page 74).

New Functions

POSIX usage brings with it an increase in real storage requirements. For example, POSIX initialization causes about 640 non-shared pages to be referenced. A subset of these pages (130 pages, in one sample) continues to be referenced by subsequent POSIX usage. Processor and I/O requirements data are provided for a selection of frequently used POSIX functions and shell commands. A loading test demonstrates that the byte file system server can handle large numbers of concurrent requests (page 83).

DCE response time and processor usage results are provided for 24 different RPC cases on 3 different configurations. Loading tests demonstrate that VM DCE servers can handle large numbers of concurrent RPC requests (page 96).

When the new GCS TSLICE option was used to decrease the GCS time slice value from 300 milliseconds (default) to 30 milliseconds for the VTAM* machine, there was little effect on system performance (page 111).

Additional Evaluations

The PC Server 500 System/390* can support many CMS users if sufficient real storage is provided. Example results are shown for a 128MB system using the FS8F0R workload where 190 CMS users are supported with 1-second average response time (page 116).

There is some decrease in processor capacity when CMS users are run in XA mode or XC mode. Comparisons were obtained on VM/ESA 1.2.2 using the FS8F0R workload. Relative to running the CMS users in 370 mode virtual machines, ITR decreased 2.0% when the CMS users were run in XA mode and 3.2% when they were run in XC mode (page 125).

Benchmark runs with the CMS 370 Accommodation Facility (CMS370AC ON) show no measurable change in overall performance relative to CMS370AC OFF. This is for the case where there are no 370-only events that need to be simulated by CP (page 131).

The effective use of minidisk caching (MDC) for guest operating systems requires adequate real and/or expanded storage. The use of MDC on a storage constrained system can result in reduced guest performance unless appropriate tuning actions are taken (page 135).

Measurements indicate that the performance of RSCS 3.2 is equivalent to RSCS 3.1. The performance of the new TCPNJE line driver is similar to the SNANJE line driver. The TCP/IP server, however, uses more processor time than the VTAM server (page 140).

For DirMaint 1.5, the DirMaint server has been rewritten in REXX and many new DirMaint functions have been added. As a result, DirMaint processor usage has increased substantially from DirMaint 1.4. Total system impact for the measured CMS-intensive environment was -4.1% ITR and -3.2% ITR with the use of compiled REXX for the DirMaint execs (page 145).

Measurements using the FS8F0R CMS-intensive workload show a 0.6% ITR improvement when migrating from VTAM 3.4.1 to VTAM 4.2.0 (page 153).

Changes That Affect Performance

This chapter contains descriptions of various changes to VM/ESA 2.1.0 that affect performance. This information is equivalent to the information on VM/ESA 2.1.0 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 9. However, some have the potential to adversely affect performance. These are listed under "Performance Considerations" on page 16. 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 17 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.

- CP
  - CP Module Linkage Changes
  - Improved VMCF Interrupt Processing
  - Virtual Channel to Channel Locking
  - CP Trace Table Default
  - MDC Fair Share Limit
  - Extended Record Cache Support
  - Improved LOCATEVM Command
- CMS
  - Compiled REXX for CMS
  - CMS Nucleus Restructure
  - CMS Pipeline Stages in Assembler
  - Data Compression Support
  - SUPERSET XEDIT Subcommand
- Other
  - GCS PAGEX Support
  - GCS SET TSLICE Command
  - VMSES/E Improvements

CP Module Linkage Changes

In order to reduce overhead in the linkage of CP modules, changes were made in three areas for some frequent linkage cases. The first was to change some dynamic linkages to use fast dynamic linkage. Fast dynamic linkage was first introduced in VM/ESA 1.1.1, and is a more efficient method to do CP dynamic linkage.

The second area was to move some highly used modules from the pageable list to the resident list to save overhead. This increases the size of the resident nucleus slightly.

The last area was to avoid writing trace entries into the trace table for calls from HCPALLVM. This module is passed the address of a routine and then calls the passed routine for every VMDBK on the system. This is done for such functions as CP monitor high frequency user state sampling. By avoiding the trace entries, much overhead is saved and the trace table is not cluttered with less useful trace entries.

Improved VMCF Interrupt Processing

The amount of processing time required to handle VMCF interrupts has been reduced. The amount of improvement is directly related to the number of pending VMCF interrupts. If there are only a few pending interrupts, the effect is negligible.

When virtual multiple processor support for VMCF was implemented in VM/ESA 1.2.1, a large system effect in processing pending VMCF interrupts was introduced. This large system effect included storage being obtained and not released. VM APAR VM58414 was written to address the storage problem. This APAR also results in a slight processor usage reduction. In VM/ESA 2.1.0, the VMCF interrupt processing was redesigned to eliminate the large system effect.

Figure 1 shows measurements of processor time used to process different numbers of pending VMCF interrupts in four environments. This figure shows the large system effects that were introduced in VM/ESA 1.2.1 and somewhat improved with APAR VM58414. The figure also shows that with the changes in VM/ESA 2.1.0, the large system effect is gone and processor usage is very similar to the VM/ESA 1.2.0 environment. Therefore, this improvement only applies to customers who are migrating from VM/ESA 1.2.1 or VM/ESA 1.2.2.

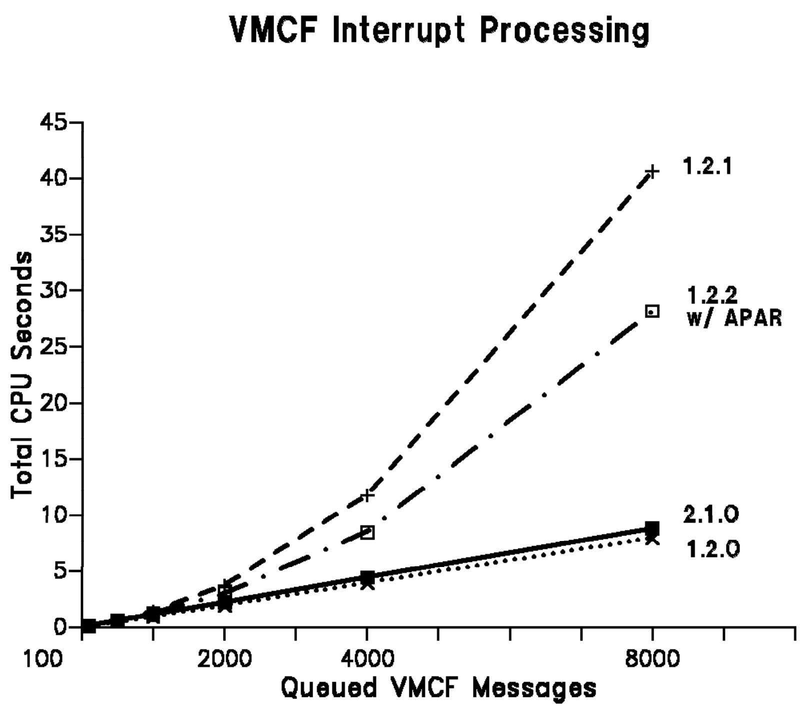

Figure 1. VMCF Interrupt Processing. Processor usage for various numbers of pending VMCF requests on a 9021-900

Virtual Channel to Channel Locking

The virtual channel to channel (VCTC) locking scheme has been improved. The new scheme uses a separate lock word for each VCTC instead of a global lock. This may improve VCTC throughput, especially in cases where a virtual machine is doing a high rate of requests through several VCTCs.

The global lock is still required in certain cases such as processing the COUPLE command.

CP Trace Table Default

The default calculation for the size of the CP trace table has been changed. This will ordinarily result in a smaller trace table for systems that take the default. As before, the trace table size can also be set explicitly.

In the past, the trace table size defaulted to 1/64th of available storage for each processor. With VM/ESA 2.1.0, that calculation is done for the master processor. However, if the result is greater than 100 pages, the master trace table size is set to 100. The trace table size for each alternate processor is then set to 75% of the size of the master processor trace table. Example savings for a 512MB 6-way system is over 46MB.

Most people feel that the new calculation represents a much better tradeoff between serviceability and real storage usage, especially considering that VM/ESA reliability has improved dramatically over the last several releases.

MDC Fair Share Limit

This improvement applies only to migrations from VM/ESA 1.2.2. It is also available on VM/ESA 1.2.2 as APAR VM59590.

Minidisk caching (MDC) has a fair share algorithm to prevent any one user from flooding the cache with data. This algorithm can be disabled with the NOMDCFS directory option. The fair share insert limit is dynamic but has a floor (minimum value). Analysis of benchmarks and customer data on VM/ESA 1.2.2 systems showed that the floor was too low and that this was degrading system performance. Accordingly, the old floor of 8 inserts per minute was increased to 150 inserts per minute.

The I/Os excluded due to the fair share limit being exceeded and the current fair share limit are reported by RTM (MDCACHE screen) and VMPRF (MINIDISK_CACHE_BY_TIME report). If, on your current system, this information shows that a significant number of I/Os are being excluded from minidisk caching due to the fair share limit being exceeded, the system may benefit from this change.

Extended Record Cache Support

Prior to VM/ESA 2.1.0, VM/ESA supported guest-use of the Record Cache I function in the 3990-6 storage control. This was for guest operating systems that build their own channel programs. Record Cache I support is now extended to users of DIAGNOSE X′A4′, DIAGNOSE X′250′, and the block-I/O facility under certain conditions. When VM knows that data being written meets the control unit′s criteria for ″regular data format″, VM sets a channel-program indicator to achieve improved performance for I/O requests issued through these I/O-service facilities. I/O requests must meet all of the following conditions:

- The request is to write data.
- The request is eligible to be stored in VM′s minidisk cache.

- The request is for a track that was previously read and found to be in standard format.[2]

The record-cache function requires the DASD Fast Write (DFW) function to be enabled and adds to the performance benefits of this function. With DFW, the control unit completes the I/O request almost immediately. The host need not wait for the data to be written (destaged) to the DASD volume, since the data are protected from loss by residing in the control unit′s nonvolatile storage (NVS) while waiting to be destaged. However, if the record is not already in the control unit′s cache when an update is received from the host, the entire track must be staged (read) into the cache before the I/O request can complete. The additional performance advantage of the record-cache function is that staging of the track containing the record can be avoided. Because VM tells the control unit that the record being written is in a standard format, the control unit knows that the record will fit within the existing format of the track when the record is ultimately destaged from the NVS.

Improved LOCATEVM Command

The LOCATEVM CP command (class G) can use a very large amount of processor time when a large search range is given. Because of this, use of this command can adversely impact system performance. As a precaution, some installations have chosen to reassign this command to a more restrictive class.

In VM/ESA 2.1.0, LOCATEVM processor requirements have been substantially reduced (by 75% in one test). While LOCATEVM can still use a significant amount of processor and paging resources, it is now less risky to leave it available to class G users.

Compiled REXX for CMS

Most of the CMS REXX execs and XEDIT macros on the S-disk are now shipped as compiled REXX files. This includes all files (except SYSPROF EXEC) that are in the CMSINST shared segment and a number of others. They make use of a subset REXX run-time library that is shipped with VM/ESA 2.1.0.

Note: Some of the CMS execs (most notably FILELIST, RDRLIST, and PEEK) are written in EXEC2. They remain in EXEC2 and are not affected by this change.

This change can significantly improve the performance of CMS intensive workloads that use REXX-implemented CMS functions such as DIRLIST, DISCARD, NAMES, NOTE, RECEIVE, SENDFILE, TELL, and VMLINK, as well as XEDIT macros such as ALL and SPLTJOIN. Processor capacity improvements exceeding 6% have been observed.

The uncompiled source files are provided on the S-disk for customers who wish to make modifications. Customers with the REXX compiler are advised to recompile the updated files before placing them back onto CMSINST so as to retain the performance advantages.

---

[2] Minidisk caching considers a track to be in standard format when it meets certain criteria. For example, all DASD records on the track must have the same length and an integral number of these records must fit into a 4KB page. There are other criteria as well; for a complete definition of standard format, see the minidisk caching chapter in *VM/ESA CP Diagnosis Reference*.

CMS Nucleus Restructure

The CMS nucleus was restructured in VM/ESA 2.1.0. This improved performance in a number of ways:

- More of the CMS nucleus has been moved above the 16MB line. This can improve performance by allowing more use of SAVEFD and by allowing more shared segments to be created that require space below the 16MB line.
  - The CMS shared system now starts at X′F00000′ instead of X′E00000′.
  - NLS language repository segments can now reside above the 16MB line.
- In prior releases, the default installation procedure placed the VMLIB, VMMTLIB, and PIPES segments below the 16MB line (even though they can be run above the line) because of the possibility that CMS could be run in a 370 mode machine. With VM/ESA 2.1.0, default installation puts these segments above the 16MB line. (VMMTLIB has been integrated into the portion of the CMS saved system that is above the line.) This change frees storage that was previously taken below the 16MB line.
- CMS now allocates some of its control blocks (such as the IUCV path table) above the 16MB line when such space is available.
- The 370-mode code has been removed from the mainline paths in CMS.
- The fast path through the SVC interrupt handler (DMSITS) has been further optimized.
- The following modules have been moved from the S-disk back into the CMS nucleus:

  DMSQRC - query COMDIR
  DMSQRE - query ENROLL
  DMSQRF - query CMS (window manager)
  DMSQRG - query CMS (window manager)
  DMSQRH - query CMS (window manager)
  DMSQRN - query NAMEDEF
  DMSQRP - query FILEPOOL
  DMSQRQ - query LIMITS, FILEWAIT, RECALL
  DMSQRT - query AUTOREAD, CMSTYPE, and so forth
  DMSQRU - query FILEDEF, LABELDEF
  DMSQRV - query INPUT, OUTPUT, SYNONYM
  DMSQRW - query libraries (MACLIB, and so forth)
  DMSQRX - query DOS, DOSPART, UPSI, DLBL
  DMSSEC - set COMDIR
  DMSSEF - set CMS (window manager)
  DMSSML - set/query MACLSUBS

  This can benefit the performance of workloads that use these functions if they had not previously been used from a shared segment.

  Note: In VM/ESA 1.2.2, these modules resided in the CMSQRYL and CMSQRYH logical segments. These segments no longer exist.

CMS Pipeline Stages in Assembler

CMS Pipelines now provide assembler macros that perform basic pipeline functions and are the building blocks for writing assembler stage commands. User-written assembler stage commands provide increased performance over similar stage commands written in REXX.

Data Compression Support

VM/ESA 2.1.0 includes data compression API support so vendors and customers can more easily create applications that exploit the use of compression services. Both a macro interface (CSRCMPSC) and a CSL interface (DMSCPR) are provided. Use of this support can save DASD space, tape storage space, and transmission line costs. The increase in processing time associated with data compression and expansion is greatly reduced on processors that have hardware compression (CMPSC instruction).

In addition, CMS and GCS support the VSE/VSAM Version 6 Release 1.0 interface for data compression. Using the COMPRESS parameter of the DEFINE function causes VSAM to automatically expand or compress data during a VSAM read or write operation, respectively. When available on the processor, the CMPSC instruction is used for this purpose. CMS and GCS system users can read and write to VSAM files that have been compressed under the control of the VSE/VSAM program. No application program changes are necessary.

SUPERSET XEDIT Subcommand

This new XEDIT subcommand performs the same function as the existing SET subcommand. However, it can be used to set multiple options in one invocation. The following CMS Productivity Aids were changed to use this subcommand: FILELIST, RDRLIST, NOTE, SENDFILE, PEEK, DIRLIST, and the EXECUTE XEDIT macro. It can also be used to improve the performance of user-written applications that include performance-sensitive XEDIT macros.

GCS PAGEX Support

You can now make use of the CP PAGEX facility with GCS. PAGEX is specified on a virtual machine basis. When PAGEX is ON and a given GCS task takes a page fault, GCS will dispatch other active GCS tasks in the virtual machine while waiting for that page fault to be resolved. This can result in increased capacity for that virtual machine to do work.

PAGEX is especially useful in cases where a virtual machine has a large number of GCS tasks and these tasks are active on an intermittent basis. A good example would be an RSCS machine with many line drivers.

If this is not the case, SET RESERVE remains the best method to minimize the effects of paging. SET RESERVE works best when the virtual machine′s reference pattern has good locality of reference and its working set size does not change much over time. In intermediate cases, the best tuning solution might be to use a combination of PAGEX ON and SET RESERVE. SET RESERVE would be used to protect the most frequently used pages, while PAGEX ON would be used to keep those page faults that do occur from serializing the whole virtual machine.

Note: PAGEX is not recommended for the VTAM machine because most VTAM execution is on one GCS task.

GCS SET TSLICE Command

In prior releases, the GCS time slice was fixed at 300 milliseconds. With VM/ESA 2.1.0, 300 milliseconds is retained as the default setting but this can be altered for any given virtual machine in a GCS group by using the new SET TSLICE GCS command.

A smaller time slice setting can be used to help avoid time-out situations when multiple tasks are involved. You can estimate whether the default time slice setting is likely to result in a time-out situation. For example, if the QUERY TSLICE command shows 100 active tasks, the maximum delay before a given task is run is 100 times 0.300, or 30 seconds. If this is more than the line time-out limit, you should set the time slice lower.

Note: Setting the time slice lower than it needs to be will tend to increase GCS dispatching overhead.

VMSES/E Improvements

The performance of the VMFBLD function has been improved. Elapsed time and processor time reductions exceeding 20% have been observed. This improvement was first introduced through VM/ESA 1.2.2 APAR VM57938.

The performance of VMFCOPY has been improved by providing an SPRODID option. In prior releases, all files that met the *fn ft fm* criteria were copied regardless of what product they belonged to. When you specify the SPRODID option, only those files that belong to the specified product are copied.

The automation of more service processing in VMSES/E 2.1.0 eliminates certain manual tasks. Therefore, the overall time required to do these tasks will decrease. See “VMSES/E” on page 70 for a list of tasks that have been automated by VMSES/E 2.1.0.

Performance Considerations

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

- Additional CMS Paging Space Requirements
- CMS Working Set Size Increase
- Potential Overlap of CMS with Shared Segments

Additional CMS Paging Space Requirements

The number of pages that are referenced during IPL CMS but are (typically) unused thereafter has increased by about 12. This increases DASD paging space requirements to some extent. Since these referenced pages must ultimately be paged out, they can also reduce performance in situations where large numbers of CMS users are logging on over a short period of time.

Many additional virtual pages in the user′s virtual machine are referenced when the POSIX environment is initialized. This occurs implicitly when the first POSIX request is made. Nearly all of these additional pages are no longer referenced if there are no subsequent POSIX requests. However, these pages will add to the number of occupied page slots on DASD. This leads to the following two recommendations:

1. If many users are (even occasionally) using the POSIX environment, take a look at whether the system′s page space is still sufficient.

2. Do not put POSIX-oriented commands such as OPENVM MOUNT in your PROFILE EXEC unless you will normally be using POSIX functions subsequent to starting CMS.

CMS Working Set Size Increase

CMS references more non-shared pages than it did in VM/ESA 1.2.2. This will tend to increase paging, especially in storage-constrained environments with large numbers of CMS users. For the CMS-intensive FS8F measurements reported in “CMS-Intensive” on page 25, working set increases ranging from 1% to 10% were observed.

The main reason for this increase is that CMS Pipelines now does a CMS multitasking call as part of its initialization. This means that users who do not use pipelines or who are already running CMS multitasking will not experience the pipelines-related working set increase.

Potential Overlap of CMS with Shared Segments

In VM/ESA 1.2.2, the CMS saved system occupied megabytes E, F, and 10. In VM/ESA 2.1.0, it occupies megabytes F, 10, 11, and 12. If your installation has defined any shared segments in megabytes 11 or 12, they will need to be moved in order to avoid overlapping CMS.

Performance Management

These changes affect the performance management of VM/ESA.

- Monitor Enhancements
- Dynamic Allocation of Subchannel Measurement Blocks
- SET THROTTLE Command
- QUERY FILEPOOL Command Extensions for BFS
- Accounting Data
- VM Performance Products

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 2.1.0.

- User State Sampling

  A number of changes were made to improve the usefulness of the user state sampling data.

  – Users doing diagnose I/O used to show up as being in simulation wait. They now appear in I/O wait.

  – Users in CP SLEEP or CP READ used to be shown as being in console function mode wait. They now appear as idle.

  – A new state, active page wait, has been added for virtual machines that have a page request outstanding but can handle it with PAGEX or asynchronous page fault handling.

- CP Configurability II

  The CP Configurability II support allows I/O devices to be added or removed from the I/O hardware configuration while VM/ESA is running. In order to track these changes, several new event I/O domain records were added, such as the delete device record (domain 6 record 15, D6/R15).

  A measurement block, sometimes referred to as a subchannel measurement block, is a control block that is associated with a given device. It contains measurement data for the device such as the I/O rate and timing values for the various components of service time. The hardware is responsible for updating this information. From the measurement block information, performance products can compute the device′s service time, I/O rate, and utilization. With the CP Configurability II support, it is now possible for a given device to not have an associated measurement block. Accordingly, information has been added to the monitor to indicate when this is the case.

  The new SET SCMEASURE command allows an administrator to enable or disable the collection of subchannel measurement data for a specific device or range of devices. An event record is created each time the SET SCMEASURE command is used.

- SET THROTTLE

  Monitor fields have been added in support of the new SET THROTTLE command. This includes:

  - whether a device has been throttled
  - the throttle rate for a device
  - the number of times I/O was delayed on a given device
  - the number of times a given user had I/O delayed due to throttle

- RAMAC* Support

  Monitor support for RAMAC is available for VM/ESA 1.2.1 and VM/ESA 1.2.2 through development APAR VM59200. This support is integrated into VM/ESA 2.1.0. Since RAMAC DASD appear to VM as either 3380s or 3390s, additional fields have been added to the device configuration data record (D1/R6) and the vary on device record (D6/R1) to indicate the actual DASD and control unit type where possible. Cache activity data records (D6/R4) have been made available for the RAMAC subsystem.

- SFS APPLDATA

  The APPLDATA domain monitor data contributed by SFS filepool servers has been extended to include counts and timings that pertain to the byte file system. These include byte file request counts for each type of request, lock conflict counts for each type of byte file lock conflict, and token callback information.

- CMS Multitasking APPLDATA

  CMS multitasking can contribute application data to the monitor in the APPLDATA (10) monitor domain. This includes the following information:

  - Thread creation and deletion counts and timings
  - Thread switch rates
  - Number of blocked threads
  - Highest number of threads and POSIX processes in use

Dynamic Allocation of Subchannel Measurement Blocks

The I/O service times and related information for a device are computed from data found in its associated subchannel measurement block, which the hardware is responsible for updating. With the new functions provided by CP Configurability II, there can now be scenarios where there is not a subchannel measurement block associated with a device. In such cases, the service times and related data are not available and are shown as zeros in the monitor data.

SET THROTTLE Command

This is a new CP command that can be used to set a maximum rate at which the system′s virtual machines, in aggregate, are permitted to initiate I/Os to a given device. This limit does not apply to I/Os initiated by CP. CP converts the specified rate into an interval representing the minimum time that must pass after one I/O is started before the next I/O to that device can start. If CP receives an I/O request to a device that has been limited by SET THROTTLE, that

I/O request is delayed, if necessary, until the minimum time interval has completed.

In multi-system configurations which have shared channels, control units, or devices, SET THROTTLE can be used to help prevent any one system from overutilizing the shared resources.

QUERY FILEPOOL Command Extensions for BFS

Information has been added to the QUERY FILEPOOL commands to provide byte file system performance information. In particular, byte file system counts and timings have been added to QUERY FILEPOOL REPORT and its subset, QUERY FILEPOOL COUNTER.

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 2.1.0. 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 2.1.0. Exception: CMS intensive workloads that make significant use of DIRLIST, DISCARD, NAMES, NOTE, RECEIVE, SENDFILE, TELL, and VMLINK, and/or XEDIT macros such as ALL and SPLTJOIN. Such workloads can experience a significant reduction in total processor time arising from CMS′s use of compiled REXX. Most of this decrease will be virtual processor time.

Milliseconds of virtual processor time (37-40)

This is the virtual time charged to a user. See the above discussion of total processor time.

Requested virtual nonspooled I/O starts (49-52)

This is a total count of requested starts. All requests may not complete. The value of this field could change, depending on the system I/O characteristics, because of two changes made to CP:

- In previous releases, this counter was incremented for each real I/O done. This included the scenario where CP splits a virtual I/O into a separate real I/O for each cylinder involved. In VM/ESA 2.1.0, this counter will be incremented only once per virtual I/O.
- In the past, virtual I/Os eligible for minidisk caching that experienced a cache miss were not always being counted. This has been corrected.

Completed virtual nonspooled I/O starts (73-76)

This is a total count of completed requests. The previous discussion of “requested virtual nonspooled I/O starts” also applies to this field.

VM Performance Products

VM Performance Reporting Facility 1.2.1 (VMPRF) requires APAR VM59656 (PTF UM27312) to run on VM/ESA 2.1.0. VMPRF at this service level includes the following functional enhancements:

- 3990-6 cache controller support
- RAMAC support
- LE/370 support
- page active and limit list data have been added to the state sample reports

Realtime Monitor VM/ESA 1.5.2 (RTM/ESA) requires APAR GC05374 (PTF UG03792) to run on VM/ESA 2.1.0. It can be run on CMS11 (or earlier) in 370 mode or on CMS12 in XA mode with CMS370AC on. If it was built on CMS12 and is run on CMS12, it will set CMS370AC on. RTM/ESA at this service level can be built using HLASM Release 2. Support for RAMAC DASD has been added.

Performance Analysis Facility/VM 1.1.3 (VMPAF) will run on VM/ESA 2.1.0 with the same support as VM/ESA 1.2.2.

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

The following processors were measured.

- 9121-742
- 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
- PC Server 500 System/390

Software

Unless otherwise noted, a pre-general-availability level of VM/ESA 2.1.0 was used for the measurements in this report. Not all of the VM/ESA 2.1.0 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 2.1.0 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 measured for this report. VM/ESA 1.2.2 was at the GA+first-RSU level (General Availability, Recommended Service Upgrade tape). The service that was part of VM/ESA 1.2.2 after the first RSU level and integrated into VM/ESA 2.1.0 can account for some of the difference between VM/ESA 1.2.2 and VM/ESA 2.1.0.

See the appropriate workload section in Appendix A, “Workloads” on page 157 for the other licensed programs′ software levels.

Format Description

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

For each group of measurements there are five sections:

1. Workload: This specifies the name of the workload associated with the measurement. For more detail on the workload, see Appendix A, “Workloads” on page 157.
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. An ″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 24.

   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.

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.

RTM Real Time Monitor, records and reports performance data for VM systems.

SPM/2 System Performance Monitor/2 provides performance data for an OS/2* system.

STARS System Trace Analysis Reports, provides various reports based on the analysis of instruction trace data.

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.

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

Migration from VM/ESA 1.2.2

This chapter explores the performance effects of migrating from VM/ESA 1.2.2 to VM/ESA 2.1.0. The following environments were measured: CMS-intensive, VSE guest, and VMSES/E.

CMS-Intensive

VM/ESA 2.1.0 has improved internal throughput rates and response times for the CMS-intensive environments measured. The ITR improvements resulting from decreased processor use can be attributed mostly to the use of compiled REXX execs and compiled XEDIT macros from the CMS system disk (a non-compiled comparison can be found in “9121-480 / Minidisk” on page 31). The following enhancements also contributed:

- CP Module Linkage Changes
- CMS Nucleus Restructure
- SUPERSET XEDIT Subcommand

For more information on these and other performance-related enhancements in VM/ESA 2.1.0, see “Changes That Affect Performance” on page 8.

The internal throughput rates and response times for these measurements are shown in Figure 2 and Figure 3 on page 26.

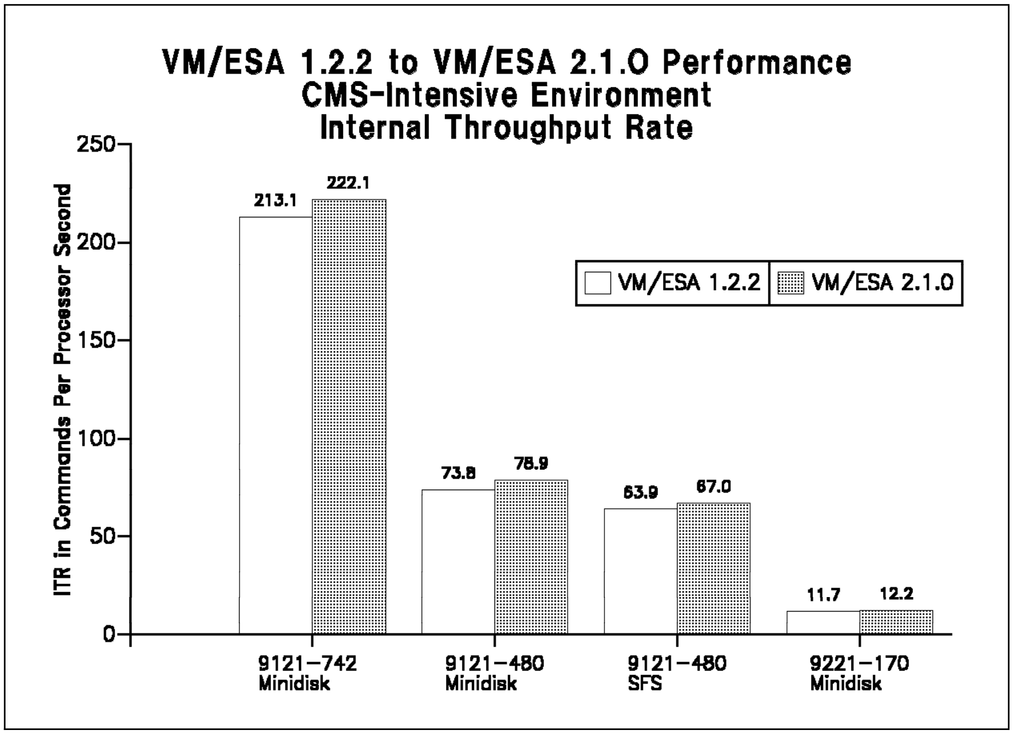

Figure 2. Internal throughput rate for the various CMS-intensive environments

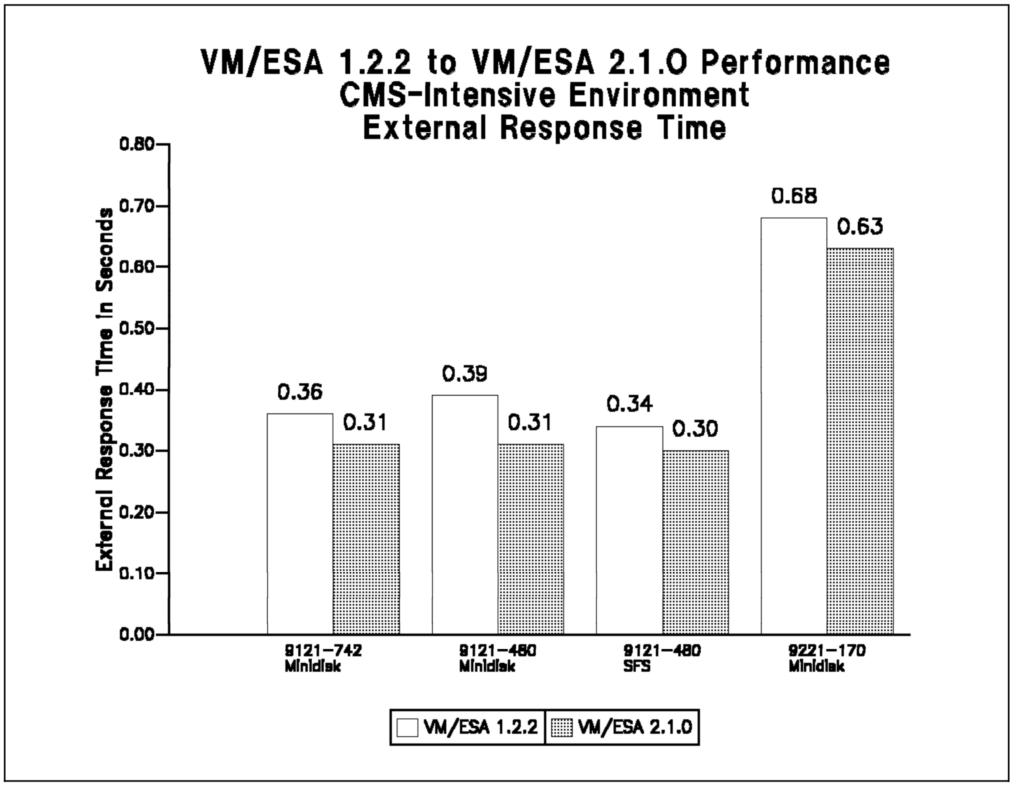

Figure 3. External response time for the various CMS-intensive environments

9121-742 / Minidisk

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-742
Processors used:	4
Storage:	
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: Several performance enhancements have been made in VM/ESA 2.1.0. The use of compiled REXX by CMS has provided most of the performance gain shown in the measurements within this section. The workload used (FS8F0R) includes a variety of these execs. For a comparison without the effects of the REXX compiler see “9121-480 / Minidisk” on page 31.

There has been some pathlength growth within GCS, causing increases in the processor usage within the VTAM machine. This has mostly been caused by inclusion of APARs to GCS.

There is a working set growth within the CMS virtual machine due to the addition of new function that was not offset by performance improvements. See "CMS Working Set Size Increase" on page 16 for more information. The increased CP processor usage is mostly caused by an increase in paging, produced by this working set growth.

Note that these measurements were made with an MDC BIAS value of 0.1 for expanded storage (using the SET MDCACHE command). Studies previously made have determined that this improves overall system performance. For more information refer to the *VM/ESA Release 2.2 Performance Report*.

The following table shows that VM/ESA 2.1.0 compared to VM/ESA 1.2.2 has improved its overall 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 13.3% and the internal throughput improved by 4.2%.

Table 1 (Page 1 of 3). Minidisk-only CMS-intensive migration from VM/ESA 1.2.2 on the 9121-742

Release **Run ID**	**1.2.2** **S47E550D**	**2.1.0** **S48E5500**	**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	0.107	0.112	0.005	4.67%
NONTRIV INT	0.350	0.335	-0.015	-4.29%
TOT INT	0.250	0.251	0.001	0.40%
TOT INT ADJ	0.247	0.234	-0.012	-5.05%
AVG FIRST (T)	0.276	0.236	-0.040	-14.51%
AVG LAST (T)	0.361	0.313	-0.048	-13.28%
Throughput				
AVG THINK (T)	26.11	26.09	-0.02	-0.06%
ETR	190.39	180.15	-10.24	-5.38%
ETR (T)	192.91	193.00	0.09	0.05%
ETR RATIO	0.987	0.933	-0.054	-5.42%
ITR (H)	213.12	222.12	9.00	4.22%
ITR	52.61	51.89	-0.72	-1.36%
EMUL ITR	78.83	82.02	3.19	4.05%
ITRR (H)	1.000	1.042	0.042	4.22%
ITRR	1.000	0.986	-0.014	-1.36%
Proc. Usage				
PBT/CMD (H)	18.769	18.009	-0.760	-4.05%
PBT/CMD	18.765	17.979	-0.786	-4.19%
CP/CMD (H)	6.659	7.037	0.378	5.67%
CP/CMD	6.220	6.580	0.360	5.78%
EMUL/CMD (H)	12.110	10.972	-1.138	-9.40%
EMUL/CMD	12.545	11.399	-1.146	-9.13%

Table 1 (Page 2 of 3). Minidisk-only CMS-intensive migration from VM/ESA 1.2.2 on the 9121-742

Release **Run ID**	**1.2.2** **S47E550D**	**2.1.0** **S48E5500**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5500**	**5500**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Processor Util.				
TOTAL (H)	362.08	347.57	-14.50	-4.01%
TOTAL	362.00	347.00	-15.00	-4.14%
UTIL/PROC (H)	90.52	86.89	-3.63	-4.01%
UTIL/PROC	90.50	86.75	-3.75	-4.14%
TOTAL EMUL (H)	233.62	211.76	-21.86	-9.36%
TOTAL EMUL	242.00	220.00	-22.00	-9.09%
MASTER TOTAL (H)	92.68	88.94	-3.73	-4.03%
MASTER TOTAL	93.00	89.00	-4.00	-4.30%
MASTER EMUL (H)	39.69	34.40	-5.28	-13.31%
MASTER EMUL	41.00	36.00	-5.00	-12.20%
TVR(H)	1.55	1.64	0.09	5.90%
TVR	1.50	1.58	0.08	5.44%
Storage				
NUCLEUS SIZE (V)	2572KB	2756KB	184KB	7.15%
TRACE TABLE (V)	800KB	800KB	0KB	0.00%
WKSET (V)	73	80	7	9.59%
PGBLPGS	233K	232K	-1K	-0.43%
PGBLPGS/USER	42.4	42.2	-0.2	-0.43%
FREEPGS	15617	16192	575	3.68%
FREE UTIL	0.92	0.92	0.00	-0.11%
SHRPGS	1799	1885	86	4.78%
Paging				
READS/SEC	643	1016	373	58.01%
WRITES/SEC	480	760	280	58.33%
PAGE/CMD	5.821	9.202	3.381	58.07%
PAGE IO RATE (V)	169.300	308.800	139.500	82.40%
PAGE IO/CMD (V)	0.878	1.600	0.722	82.31%
XSTOR IN/SEC	823	595	-228	-27.70%
XSTOR OUT/SEC	1429	1558	129	9.03%
XSTOR/CMD	11.674	11.155	-0.518	-4.44%
FAST CLR/CMD	8.994	8.601	-0.393	-4.37%
Queues				
DISPATCH LIST	101.18	101.50	0.32	0.31%
ELIGIBLE LIST	0.02	0.00	-0.02	-100.00%
I/O				
VIO RATE	1794	1847	53	2.95%
VIO/CMD	9.300	9.570	0.270	2.91%
RIO RATE (V)	544	693	149	27.39%
RIO/CMD (V)	2.820	3.591	0.771	27.33%
NONPAGE RIO/CMD (V)	1.942	1.991	0.048	2.49%
DASD RESP TIME (V)	19.900	20.200	0.300	1.51%
MDC REAL SIZE (MB)	33.2	30.2	-3.0	-8.97%
MDC XSTOR SIZE (MB)	63.5	63.6	0.1	0.10%
MDC READS (I/Os)	552	596	44	7.97%
MDC WRITES (I/Os)	26	26	0	0.00%
MDC AVOID	512	556	44	8.59%
MDC HIT RATIO	0.92	0.93	0.01	1.09%

Table 1 (Page 3 of 3). Minidisk-only CMS-intensive migration from VM/ESA 1.2.2 on the 9121-742

Release Run ID	1.2.2 S47E550D	2.1.0 S48E5500	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	20.381	20.703	0.322	1.58%
DIAG/CMD	25.628	23.968	-1.660	-6.48%
DIAG 04/CMD	0.917	0.933	0.016	1.70%
DIAG 08/CMD	0.749	0.750	0.001	0.13%
DIAG 0C/CMD	1.126	1.125	0.000	-0.02%
DIAG 14/CMD	0.024	0.025	0.000	0.18%
DIAG 58/CMD	1.248	1.248	0.000	0.01%
DIAG 98/CMD	0.324	0.382	0.058	17.90%
DIAG A4/CMD	3.571	3.769	0.198	5.54%
DIAG A8/CMD	2.814	2.826	0.011	0.40%
DIAG 214/CMD	13.708	11.680	-2.028	-14.79%
SIE/CMD	57.021	56.994	-0.027	-0.05%
SIE INTCPT/CMD	37.634	38.756	1.122	2.98%
FREE TOTL/CMD	44.792	45.165	0.372	0.83%
VTAM Machines				
WKSET (V)	4140	4140	0	0.00%
TOT CPU/CMD (V)	2.7502	3.0437	0.2935	10.67%
CP CPU/CMD (V)	1.2268	1.3147	0.0879	7.16%
VIRT CPU/CMD (V)	1.5234	1.7290	0.2056	13.50%
DIAG 98/CMD (V)	0.324	0.381	0.057	17.63%

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:	
Real:	256MB (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	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	560	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1900	Users	3MB/XC	100		

Measurement Discussion: As in the previous section, the dominating performance improvement is the compiled CMS system REXX execs and XEDIT macros contained on the S-disk. Also shown later in this section is a comparison to VM/ESA 1.2.2 without the use of these compiled execs.

The following table shows that VM/ESA 2.1.0 compared to VM/ESA 1.2.2 has improved its overall 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 20.5% and the internal throughput improved by 7.0%.

Migration: CMS-Intensive

Table 2 (Page 1 of 2). Minidisk-only CMS-intensive migration from VM/ESA 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27E1909**	**2.1.0** **L28E190M**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.130	0.123	-0.007	-5.38%
NONTRIV INT	0.456	0.367	-0.089	-19.52%
TOT INT	0.345	0.286	-0.059	-17.10%
TOT INT ADJ	0.307	0.251	-0.056	-18.17%
AVG FIRST (T)	0.272	0.224	-0.048	-17.65%
AVG LAST (T)	0.388	0.309	-0.080	-20.46%
Throughput				
AVG THINK (T)	26.15	26.16	0.01	0.06%
ETR	59.34	58.75	-0.59	-0.99%
ETR (T)	66.67	66.87	0.20	0.30%
ETR RATIO	0.890	0.879	-0.011	-1.29%
ITR (H)	73.76	78.93	5.17	7.00%
ITR	32.85	34.69	1.84	5.61%
EMUL ITR	47.65	51.81	4.16	8.72%
ITRR (H)	1.000	1.070	0.070	7.00%
ITRR	1.000	1.056	0.056	5.61%
Proc. Usage				
PBT/CMD (H)	27.113	25.339	-1.775	-6.54%
PBT/CMD	27.147	25.272	-1.876	-6.91%
CP/CMD (H)	9.009	8.958	-0.051	-0.57%
CP/CMD	8.399	8.374	-0.025	-0.30%
EMUL/CMD (H)	18.104	16.381	-1.723	-9.52%
EMUL/CMD	18.748	16.898	-1.851	-9.87%
Processor Util.				
TOTAL (H)	180.77	169.45	-11.32	-6.26%
TOTAL	181.00	169.00	-12.00	-6.63%
UTIL/PROC (H)	90.39	84.72	-5.66	-6.26%
UTIL/PROC	90.50	84.50	-6.00	-6.63%
TOTAL EMUL (H)	120.71	109.55	-11.16	-9.25%
TOTAL EMUL	125.00	113.00	-12.00	-9.60%
MASTER TOTAL (H)	89.96	84.07	-5.90	-6.56%
MASTER TOTAL	90.00	84.00	-6.00	-6.67%
MASTER EMUL (H)	53.35	47.88	-5.47	-10.26%
MASTER EMUL	55.00	50.00	-5.00	-9.09%
TVR(H)	1.50	1.55	0.05	3.29%
TVR	1.45	1.50	0.05	3.29%
Storage				
NUCLEUS SIZE (V)	2572KB	2756KB	184KB	7.15%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	84	86	2	2.38%
PGBLPGS	55084	55088	4	0.01%
PGBLPGS/USER	29.0	29.0	0.0	0.01%
FREEPGS	5406	5570	164	3.03%
FREE UTIL	0.95	0.97	0.02	1.91%
SHRPGS	1313	1361	48	3.66%

Table 2 (Page 2 of 2). Minidisk-only CMS-intensive migration from VM/ESA 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27E1909**	**2.1.0** **L28E190M**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	606	669	63	10.40%
WRITES/SEC	445	450	5	1.12%
PAGE/CMD	15.763	16.733	0.970	6.15%
PAGE IO RATE (V)	181.700	187.900	6.200	3.41%
PAGE IO/CMD (V)	2.725	2.810	0.085	3.10%
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	8.954	8.524	-0.431	-4.81%
Queues				
DISPATCH LIST	41.98	36.88	-5.11	-12.16%
ELIGIBLE LIST	0.00	0.02	0.02	na
I/O				
VIO RATE	671	699	28	4.17%
VIO/CMD	10.064	10.453	0.389	3.86%
RIO RATE (V)	393	389	-4	-1.02%
RIO/CMD (V)	5.894	5.817	-0.077	-1.31%
NONPAGE RIO/CMD (V)	3.169	3.007	-0.162	-5.11%
DASD RESP TIME (V)	19.200	19.700	0.500	2.60%
MDC REAL SIZE (MB)	41.4	39.5	-1.9	-4.52%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	191	207	16	8.38%
MDC WRITES (I/Os)	9.55	9.49	-0.06	-0.63%
MDC AVOID	180	196	16	8.89%
MDC HIT RATIO	0.94	0.94	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.906	13.856	-0.051	-0.36%
DIAG/CMD	27.861	26.448	-1.413	-5.07%
DIAG 04/CMD	2.401	2.480	0.080	3.31%
DIAG 08/CMD	0.752	0.748	-0.003	-0.44%
DIAG 0C/CMD	1.126	1.126	0.000	-0.01%
DIAG 14/CMD	0.025	0.024	0.000	-1.10%
DIAG 58/CMD	1.248	1.247	-0.001	-0.10%
DIAG 98/CMD	1.081	1.250	0.169	15.63%
DIAG A4/CMD	3.590	3.802	0.212	5.91%
DIAG A8/CMD	2.823	2.832	0.009	0.31%
DIAG 214/CMD	13.665	11.707	-1.958	-14.33%
SIE/CMD	54.070	53.803	-0.267	-0.49%
SIE INTCPT/CMD	34.605	35.510	0.905	2.62%
FREE TOTL/CMD	49.495	50.110	0.614	1.24%
VTAM Machines				
WKSET (V)	559	551	-8	-1.43%
TOT CPU/CMD (V)	3.7996	4.1538	0.3542	9.32%
CP CPU/CMD (V)	1.4415	1.5203	0.0788	5.47%
VIRT CPU/CMD (V)	2.3581	2.6335	0.2754	11.68%
DIAG 98/CMD (V)	1.081	1.250	0.170	15.69%

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

Migration: CMS-Intensive

The following table compares VM/ESA 2.1.0 without using the compiled version of the REXX execs and XEDIT macros to VM/ESA 1.2.2. This comparison may more closely represent the performance that can be expected if your system already utilizes the REXX compiler for the REXX programs and XEDIT macros found on the CMS system disk. The key indicators of external response time (AVG LAST(T)) and internal throughput rate (ITR(H)) both show a slight degradation. The external response time increased by 2.1% and the internal throughput decreased by 1.1%.

Table 3 (Page 1 of 3). Minidisk-only CMS-intensive migration (without complied REXX) from VM/ESA 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27E1909**	**2.1.0** **L28E190K**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.130	0.136	0.006	4.62%
NONTRIV INT	0.456	0.468	0.012	2.63%
TOT INT	0.345	0.356	0.011	3.19%
TOT INT ADJ	0.307	0.317	0.010	3.18%
AVG FIRST (T)	0.272	0.280	0.008	2.94%
AVG LAST (T)	0.388	0.397	0.008	2.06%
Throughput				
AVG THINK (T)	26.15	26.19	0.04	0.15%
ETR	59.34	59.33	-0.01	-0.02%
ETR (T)	66.67	66.67	0.00	-0.01%
ETR RATIO	0.890	0.890	0.000	-0.01%
ITR (H)	73.76	72.96	-0.80	-1.09%
ITR	32.85	32.47	-0.38	-1.14%
EMUL ITR	47.65	46.82	-0.83	-1.75%
ITRR (H)	1.000	0.989	-0.011	-1.09%
ITRR	1.000	0.989	-0.011	-1.14%
Proc. Usage				
PBT/CMD (H)	27.113	27.411	0.298	1.10%
PBT/CMD	27.147	27.449	0.302	1.11%
CP/CMD (H)	9.009	9.000	-0.008	-0.09%
CP/CMD	8.399	8.400	0.001	0.01%
EMUL/CMD (H)	18.104	18.410	0.306	1.69%
EMUL/CMD	18.748	19.050	0.301	1.61%
Processor Util.				
TOTAL (H)	180.77	182.74	1.97	1.09%
TOTAL	181.00	183.00	2.00	1.10%
UTIL/PROC (H)	90.39	91.37	0.99	1.09%
UTIL/PROC	90.50	91.50	1.00	1.10%
TOTAL EMUL (H)	120.71	122.74	2.03	1.68%
TOTAL EMUL	125.00	127.00	2.00	1.60%
MASTER TOTAL (H)	89.96	91.01	1.05	1.16%
MASTER TOTAL	90.00	91.00	1.00	1.11%
MASTER EMUL (H)	53.35	54.80	1.45	2.71%
MASTER EMUL	55.00	57.00	2.00	3.64%
TVR(H)	1.50	1.49	-0.01	-0.58%
TVR	1.45	1.44	-0.01	-0.49%

Table 3 (Page 2 of 3). Minidisk-only CMS-intensive migration (without complied REXX) from VM/ESA 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27E1909**	**2.1.0** **L28E190K**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Storage				
NUCLEUS SIZE (V)	2572KB	2756KB	184KB	7.15%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	84	90	6	7.14%
PGBLPGS	55084	55050	-34	-0.06%
PGBLPGS/USER	29.0	29.0	0.0	-0.06%
FREEPGS	5406	5602	196	3.63%
FREE UTIL	0.95	0.96	0.01	1.33%
SHRPGS	1313	1319	6	0.46%
Paging				
READS/SEC	606	672	66	10.89%
WRITES/SEC	445	461	16	3.60%
PAGE/CMD	15.763	16.995	1.231	7.81%
PAGE IO RATE (V)	181.700	205.400	23.700	13.04%
PAGE IO/CMD (V)	2.725	3.081	0.356	13.05%
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	8.954	9.240	0.286	3.19%
Queues				
DISPATCH LIST	41.98	41.37	-0.61	-1.46%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	671	678	7	1.04%
VIO/CMD	10.064	10.170	0.106	1.05%
RIO RATE (V)	393	388	-5	-1.27%
RIO/CMD (V)	5.894	5.820	-0.075	-1.26%
NONPAGE RIO/CMD (V)	3.169	2.739	-0.430	-13.58%
DASD RESP TIME (V)	19.200	19.400	0.200	1.04%
MDC REAL SIZE (MB)	41.4	38.0	-3.4	-8.28%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	191	203	12	6.28%
MDC WRITES (I/Os)	9.55	9.47	-0.08	-0.84%
MDC AVOID	180	192	12	6.67%
MDC HIT RATIO	0.94	0.94	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.906	13.890	-0.016	-0.11%
DIAG/CMD	27.861	28.046	0.185	0.67%
DIAG 04/CMD	2.401	2.400	-0.001	-0.05%
DIAG 08/CMD	0.752	0.750	-0.001	-0.18%
DIAG 0C/CMD	1.126	1.126	0.000	-0.01%
DIAG 14/CMD	0.025	0.025	0.000	-0.18%
DIAG 58/CMD	1.248	1.247	-0.001	-0.09%
DIAG 98/CMD	1.081	1.044	-0.036	-3.37%
DIAG A4/CMD	3.590	3.760	0.171	4.75%
DIAG A8/CMD	2.823	2.802	-0.021	-0.75%
DIAG 214/CMD	13.665	13.742	0.077	0.56%
SIE/CMD	54.070	54.599	0.529	0.98%
SIE INTCPT/CMD	34.605	34.397	-0.207	-0.60%
FREE TOTL/CMD	49.495	49.559	0.064	0.13%

Table 3 (Page 3 of 3). Minidisk-only CMS-intensive migration (without complied REXX) from VM/ESA 1.2.2 on the 9121-480

Release Run ID	1.2.2 L27E1909	2.1.0 L28E190K	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
VTAM Machines				
WKSET (V)	559	573	14	2.50%
TOT CPU/CMD (V)	3.7996	4.0416	0.2420	6.37%
CP CPU/CMD (V)	1.4415	1.4833	0.0418	2.90%
VIRT CPU/CMD (V)	2.3581	2.5583	0.2002	8.49%
DIAG 98/CMD (V)	1.081	1.045	-0.036	-3.29%

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:	
Real:	256MB (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	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*
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		

Measurement Discussion: Internal throughput (ITR(H)) improved by 4.9%, while external response time (AVG LAST(T)) improved by 12%. These improvements were mostly due to the use of compiled REXX by CMS in VM/ESA 2.1.0.

The percentage ITR improvement is somewhat less than the 7.0% improvement observed for the corresponding minidisk-only measurements (“9121-480 / Minidisk” on page 31). The primary reason for this is that, in the SFS case, the same absolute decrease in processor time per command resulting from the compiled REXX item is divided by a larger base processor time per command, resulting in a smaller percentage decrease. In addition, there has been some increase in SFS server processor usage due to service and the byte file system support. This has offset some of the processor usage improvement from the compiled REXX item.

The 5% increase in SFS server working set is due to an in-memory SFS call trace table that was added in VM/ESA 2.1.0.

Table 4 (Page 1 of 3). SFS CMS-intensive migration from VM/ESA 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27S1625**	**2.1.0** **L28S1625**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.121	0.115	-0.006	-4.96%
NONTRIV INT	0.440	0.378	-0.062	-14.09%
TOT INT	0.333	0.291	-0.042	-12.61%
TOT INT ADJ	0.292	0.255	-0.038	-12.85%
AVG FIRST (T)	0.236	0.210	-0.026	-11.23%
AVG LAST (T)	0.341	0.299	-0.042	-12.44%
Throughput				
AVG THINK (T)	26.19	26.17	-0.02	-0.10%
ETR	50.06	49.86	-0.20	-0.40%
ETR (T)	57.07	57.00	-0.07	-0.13%
ETR RATIO	0.877	0.875	-0.002	-0.27%
ITR (H)	63.88	67.02	3.14	4.92%
ITR	28.02	29.33	1.30	4.65%
EMUL ITR	41.31	44.53	3.23	7.82%
ITRR (H)	1.000	1.049	0.049	4.92%
ITRR	1.000	1.047	0.047	4.65%
Proc. Usage				
PBT/CMD (H)	31.311	29.843	-1.467	-4.69%
PBT/CMD	31.366	29.827	-1.539	-4.91%
CP/CMD (H)	10.696	10.845	0.148	1.39%
CP/CMD	10.163	10.176	0.013	0.13%
EMUL/CMD (H)	20.614	18.998	-1.616	-7.84%
EMUL/CMD	21.203	19.651	-1.552	-7.32%
Processor Util.				
TOTAL (H)	178.68	170.09	-8.59	-4.81%
TOTAL	179.00	170.00	-9.00	-5.03%
UTIL/PROC (H)	89.34	85.05	-4.30	-4.81%
UTIL/PROC	89.50	85.00	-4.50	-5.03%
TOTAL EMUL (H)	117.64	108.28	-9.36	-7.95%
TOTAL EMUL	121.00	112.00	-9.00	-7.44%
MASTER TOTAL (H)	89.15	84.57	-4.58	-5.14%
MASTER TOTAL	89.00	85.00	-4.00	-4.49%
MASTER EMUL (H)	53.04	48.43	-4.61	-8.68%
MASTER EMUL	55.00	50.00	-5.00	-9.09%
TVR(H)	1.52	1.57	0.05	3.42%
TVR	1.48	1.52	0.04	2.60%

Table 4 (Page 2 of 3). SFS CMS-intensive migration from VM/ESA 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27S1625**	**2.1.0** **L28S1625**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Storage				
NUCLEUS SIZE (V)	2572KB	2756KB	184KB	7.15%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	81	82	1	1.23%
PGBLPGS	56281	56106	-175	-0.31%
PGBLPGS/USER	34.7	34.6	-0.1	-0.31%
FREEPGS	4741	4871	130	2.74%
FREE UTIL	0.92	0.95	0.03	3.06%
SHRPGS	1471	1542	71	4.83%
Paging				
READS/SEC	524	550	26	4.96%
WRITES/SEC	365	372	7	1.92%
PAGE/CMD	15.578	16.177	0.599	3.84%
PAGE IO RATE (V)	140.100	143.300	3.200	2.28%
PAGE IO/CMD (V)	2.455	2.514	0.059	2.41%
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	8.709	8.281	-0.428	-4.91%
Queues				
DISPATCH LIST	39.23	34.93	-4.30	-10.97%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	580	596	16	2.76%
VIO/CMD	10.163	10.457	0.294	2.89%
RIO RATE (V)	332	345	13	3.92%
RIO/CMD (V)	5.818	6.053	0.235	4.05%
NONPAGE RIO/CMD (V)	3.363	3.539	0.176	5.24%
DASD RESP TIME (V)	18.300	18.200	-0.100	-0.55%
MDC REAL SIZE (MB)	67.3	65.6	-1.8	-2.61%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	151	162	11	7.28%
MDC WRITES (I/Os)	14	14	0	0.00%
MDC AVOID	126	137	11	8.73%
MDC HIT RATIO	0.83	0.84	0.01	1.20%
PRIVOPs				
PRIVOP/CMD	20.390	20.508	0.118	0.58%
DIAG/CMD	26.113	24.363	-1.750	-6.70%
DIAG 04/CMD	2.660	2.657	-0.003	-0.10%
DIAG 08/CMD	0.748	0.751	0.003	0.40%
DIAG 0C/CMD	1.148	1.148	0.000	-0.02%
DIAG 14/CMD	0.024	0.025	0.000	1.13%
DIAG 58/CMD	1.248	1.249	0.001	0.09%
DIAG 98/CMD	1.312	1.432	0.120	9.12%
DIAG A4/CMD	1.983	2.164	0.181	9.14%
DIAG A8/CMD	2.613	2.601	-0.013	-0.49%
DIAG 214/CMD	13.240	11.112	-2.128	-16.07%
SIE/CMD	60.980	60.935	-0.045	-0.07%
SIE INTCPT/CMD	41.467	42.654	1.188	2.86%
FREE TOTL/CMD	52.604	53.268	0.663	1.26%

Table 4 (Page 3 of 3). SFS CMS-intensive migration from VM/ESA 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27S1625**	**2.1.0** **L28S1625**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
VTAM Machines				
WKSET (V)	505	502	-3	-0.59%
TOT CPU/CMD (V)	4.0011	4.3571	0.3560	8.90%
CP CPU/CMD (V)	1.5089	1.5888	0.0799	5.30%
VIRT CPU/CMD (V)	2.4922	2.7683	0.2761	11.08%
DIAG 98/CMD (V)	1.312	1.433	0.121	9.19%
SFS Servers				
WKSET (V)	3205	3364	159	4.96%
TOT CPU/CMD (V)	3.1931	3.3336	0.1405	4.40%
CP CPU/CMD (V)	1.4311	1.4621	0.0310	2.17%
VIRT CPU/CMD (V)	1.7620	1.8715	0.1095	6.21%
FP REQ/CMD(Q)	1.119	1.142	0.023	2.06%
IO/CMD (Q)	1.578	1.589	0.011	0.70%
IO TIME/CMD (Q)	0.021	0.022	0.001	4.76%
SFS TIME/CMD (Q)	0.027	0.030	0.003	11.11%

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 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27S1625**	**2.1.0** **L28S1625**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Close File Requests	0.3638	0.3617	-0.0021	-0.58%
Commit Requests	0.0163	0.0163	0.0000	0.00%
Connect Requests	0.0078	0.0078	0.0000	0.00%
Delete File Requests	0.0734	0.0732	-0.0002	-0.27%
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	0.2188	0.2167	-0.0021	-0.96%
Open File Replace Requests	0.1205	0.1207	0.0002	0.17%
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	0.0212	0.0212	0.0000	0.00%
Read File Requests	0.1451	0.1715	0.0264	18.19%
Refresh Directory Requests	0.0227	0.0227	0.0000	0.00%
Rename Requests	0.0049	0.0049	0.0000	0.00%
Unlock Requests	0.0246	0.0246	0.0000	0.00%
Write File Requests	0.0504	0.0508	0.0004	0.79%
Total File Pool Requests	1.1188	1.1416	0.0228	2.04%
File Pool Request Service Time	27.3494	30.1353	2.7859	10.19%
Local File Pool Requests	1.1188	1.1416	0.0228	2.04%
Begin LUWs	0.4434	0.4419	-0.0015	-0.34%
Agent Holding Time (msec)	89.3498	81.6177	-7.7321	-8.65%
SAC Calls	5.4329	5.4533	0.0204	0.38%
Catalog Lock Conflicts	0.0005	0.0011	0.0006	120.00%
Total Lock Conflicts	0.0005	0.0011	0.0006	120.00%
Lock Wait Time (msec)	0.0128	0.0245	0.0117	91.41%
File Blocks Read	0.9047	0.8983	-0.0064	-0.71%
File Blocks Written	0.5001	0.4968	-0.0033	-0.66%
Catalog Blocks Read	0.4953	0.4848	-0.0105	-2.12%
Catalog Blocks Written	0.2569	0.2526	-0.0043	-1.67%
Control Minidisk Blocks Written	0.0499	0.0496	-0.0003	-0.60%
Log Blocks Written	0.4569	0.4565	-0.0004	-0.09%
Total DASD Block Transfers	2.6638	2.6386	-0.0252	-0.95%
BIO Requests to Read File Block	0.3912	0.4139	0.0227	5.80%
BIO Requests to Write File Blocks	0.1792	0.1783	-0.0009	-0.50%
BIO Requests to Read Catalog Blks	0.4953	0.4848	-0.0105	-2.12%
BIO Requests to Write Catalog Blks	0.2081	0.2036	-0.0045	-2.16%
BIO Requests to Write Ctl Mdsk Blks	0.0020	0.0020	0.0000	0.00%
BIO Requests to Write Log Blocks	0.3980	0.3977	-0.0003	-0.08%
Total BIO Requests	1.6739	1.6804	0.0065	0.39%
Total BIO Request Time (msec)	21.4774	21.6018	0.1244	0.58%
I/O Requests to Read File Blocks	0.2687	0.2982	0.0295	10.98%
I/O Requests to Write File Blocks	0.1964	0.1934	-0.0030	-1.53%
I/O Requests to Read Catalog Blks	0.4953	0.4848	-0.0105	-2.12%
I/O Requests to Write Catalog Blks	0.2150	0.2104	-0.0046	-2.14%
I/O Requests to Write Ctl Mdsk Blks	0.0039	0.0039	0.0000	0.00%
I/O Requests to Write Log Blocks	0.3984	0.3980	-0.0004	-0.10%
Total I/O Requests	1.5776	1.5887	0.0111	0.70%

Table 5 (Page 2 of 2). SFS CMS-intensive migration from VM/ESA 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27S1625**	**2.1.0** **L28S1625**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Get Logname Requests	0.0032	0.0032	0.0000	0.00%
Get LUWID Requests	0.0032	0.0032	0.0000	0.00%
Total CRR Requests	0.0065	0.0065	0.0000	0.00%
CRR Request Service Time (msec)	0.0755	0.0756	0.0001	0.13%
Log I/O Requests	0.0065	0.0065	0.0000	0.00%

Note: Query Filepool Counters — normalized by command

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 1.2.2 on the 9121-480

Release **Run ID**	**1.2.2** **L27S1625**	**2.1.0** **L28S1625**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1620**	**1620**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Agents Held	5.1	4.7	-0.4	-8.77%
Agents In-call	1.6	1.7	0.2	10.05%
Avg LUW Time (msec)	201.5	184.7	-16.8	-8.34%
Avg File Pool Request Time (msec)	24.4	26.4	2.0	7.99%
Avg Lock Wait Time (msec)	25.6	22.3	-3.3	-13.00%
SAC Calls / FP Request	4.86	4.78	-0.08	-1.63%
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	781	781	na
Checkpoints Taken (delta)	32	32	0	0.00%
Checkpoint Duration (sec)	1.8	2.6	0.8	45.23%
Seconds Between Checkpoints	60.2	60.2	0.0	0.00%
Checkpoint Utilization	2.9	4.3	1.3	45.23%
BIO Request Time (msec)	12.83	12.86	0.02	0.19%
Blocking Factor (Blocks/BIO)	1.59	1.57	-0.02	-1.33%
Chaining Factor (Blocks/IO)	1.69	1.66	-0.03	-1.64%

Note: Query Filepool Counters — derived results

9221-170 / Minidisk

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	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 2.1.0 compared to VM/ESA 1.2.2 has improved its overall performance characteristics. The external response time (AVG LAST(T)) decreased by 7.2% and the internal throughput rate (ITR(H)) improved by 3.8%.

Table 7 (Page 1 of 2). Minidisk-only CMS-intensive migration from VM/ESA 1.2.2 on the 9221-170

Release **Run ID**	**1.2.2** **H17E0304**	**2.1.0** **H18E0303**	**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	0.230	0.221	-0.009	-3.91%
NONTRIV INT	0.976	0.909	-0.067	-6.86%
TOT INT	0.749	0.704	-0.045	-6.01%
TOT INT ADJ	0.639	0.599	-0.040	-6.21%
AVG FIRST (T)	0.387	0.358	-0.029	-7.49%
AVG LAST (T)	0.680	0.631	-0.049	-7.21%
Throughput				
AVG THINK (T)	28.03	28.09	0.06	0.21%
ETR	8.90	8.89	-0.01	-0.11%
ETR (T)	10.43	10.44	0.01	0.10%
ETR RATIO	0.853	0.851	-0.002	-0.21%
ITR (H)	11.72	12.17	0.45	3.84%
ITR	10.00	10.37	0.37	3.65%
EMUL ITR	14.90	15.86	0.96	6.45%
ITRR (H)	1.000	1.038	0.038	3.84%
ITRR	1.000	1.037	0.037	3.65%
Proc. Usage				
PBT/CMD (H)	85.357	82.200	-3.157	-3.70%
PBT/CMD	85.310	82.353	-2.957	-3.47%
CP/CMD (H)	33.695	33.931	0.236	0.70%
CP/CMD	27.798	28.728	0.930	3.35%
EMUL/CMD (H)	51.663	48.269	-3.393	-6.57%
EMUL/CMD	57.512	53.625	-3.887	-6.76%
Processor Util.				
TOTAL (H)	89.05	85.84	-3.21	-3.60%
TOTAL	89.00	86.00	-3.00	-3.37%
TOTAL EMUL (H)	53.90	50.41	-3.49	-6.48%
TOTAL EMUL	60.00	56.00	-4.00	-6.67%
TVR(H)	1.65	1.70	0.05	3.07%
TVR	1.48	1.54	0.05	3.53%
Storage				
NUCLEUS SIZE (V)	2556KB	2740KB	184KB	7.20%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
WKSET (V)	87	90	3	3.45%
PGBLPGS	13506	13427	-79	-0.58%
PGBLPGS/USER	45.0	44.8	-0.3	-0.58%
FREEPGS	953	972	19	1.99%
FREE UTIL	0.89	0.90	0.01	0.82%
SHRPGS	1101	1199	98	8.90%
Paging				
READS/SEC	84	86	2	2.38%
WRITES/SEC	69	70	1	1.45%
PAGE/CMD	14.666	14.939	0.273	1.86%
PAGE IO RATE (V)	26.300	27.000	0.700	2.66%
PAGE IO/CMD (V)	2.521	2.586	0.065	2.56%
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	8.819	8.427	-0.392	-4.44%

Table 7 (Page 2 of 2). Minidisk-only CMS-intensive migration from VM/ESA 1.2.2 on the 9221-170

Release **Run ID**	**1.2.2** **H17E0304**	**2.1.0** **H18E0303**	**Difference**	**%Difference**
Environment				
Real Storage	**64MB**	**64MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**300**	**300**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**1**	**1**		
Queues				
DISPATCH LIST	11.37	11.38	0.01	0.05%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	117	119	2	1.71%
VIO/CMD	11.215	11.395	0.180	1.61%
RIO RATE (V)	69	70	1	1.45%
RIO/CMD (V)	6.614	6.703	0.089	1.35%
NONPAGE RIO/CMD (V)	4.093	4.118	0.025	0.60%
DASD RESP TIME (V)	24.200	24.400	0.200	0.83%
MDC REAL SIZE (MB)	11.9	11.2	-0.7	-5.77%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	30	32	2	6.67%
MDC WRITES (I/Os)	1.51	1.49	-0.02	-1.32%
MDC AVOID	27	29	2	7.41%
MDC HIT RATIO	0.91	0.92	0.01	1.10%
PRIVOPs				
PRIVOP/CMD	14.514	14.490	-0.023	-0.16%
DIAG/CMD	32.992	31.268	-1.724	-5.23%
DIAG 04/CMD	6.268	6.316	0.048	0.77%
DIAG 08/CMD	0.747	0.747	0.000	0.00%
DIAG 0C/CMD	1.128	1.129	0.001	0.10%
DIAG 14/CMD	0.024	0.024	0.000	-0.10%
DIAG 58/CMD	1.251	1.251	0.001	0.05%
DIAG 98/CMD	2.286	2.270	-0.016	-0.69%
DIAG A4/CMD	3.587	3.795	0.208	5.79%
DIAG A8/CMD	2.821	2.827	0.006	0.23%
DIAG 214/CMD	13.729	11.669	-2.060	-15.01%
SIE/CMD	63.455	62.914	-0.541	-0.85%
SIE INTCPT/CMD	44.419	44.669	0.250	0.56%
FREE TOTL/CMD	58.759	58.988	0.229	0.39%
VTAM Machines				
WKSET (V)	288	274	-14	-4.86%
TOT CPU/CMD (V)	17.6531	18.3805	0.7274	4.12%
CP CPU/CMD (V)	6.8962	6.9958	0.0996	1.44%
VIRT CPU/CMD (V)	10.7570	11.3848	0.6278	5.84%
DIAG 98/CMD (V)	2.287	2.270	-0.017	-0.73%

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

VSE/ESA Guest

In the following three sections, VSE/ESA* 2.1.0[3] guest performance measurement results are presented and discussed for the DYNAPACE batch workload and the VSECICS transaction processing workload. These sections compare VSE/ESA guest performance on VM/ESA 1.2.2 to VM/ESA 2.1.0 and to VSE/ESA native.

9121-320 / DYNAPACE

This section examines VSE/ESA 2.1.0 guest performance of VM/ESA 1.2.2 compared to VM/ESA 2.1.0 running the DYNAPACE workload on a 9121-320. DYNAPACE is a batch-only workload and is characterized by heavy I/O. See Appendix A, "Workloads" on page 157 for a detailed description of the workload. Because the 9121-320 is a uniprocessor, the VSE/ESA Standard Dispatcher is used.

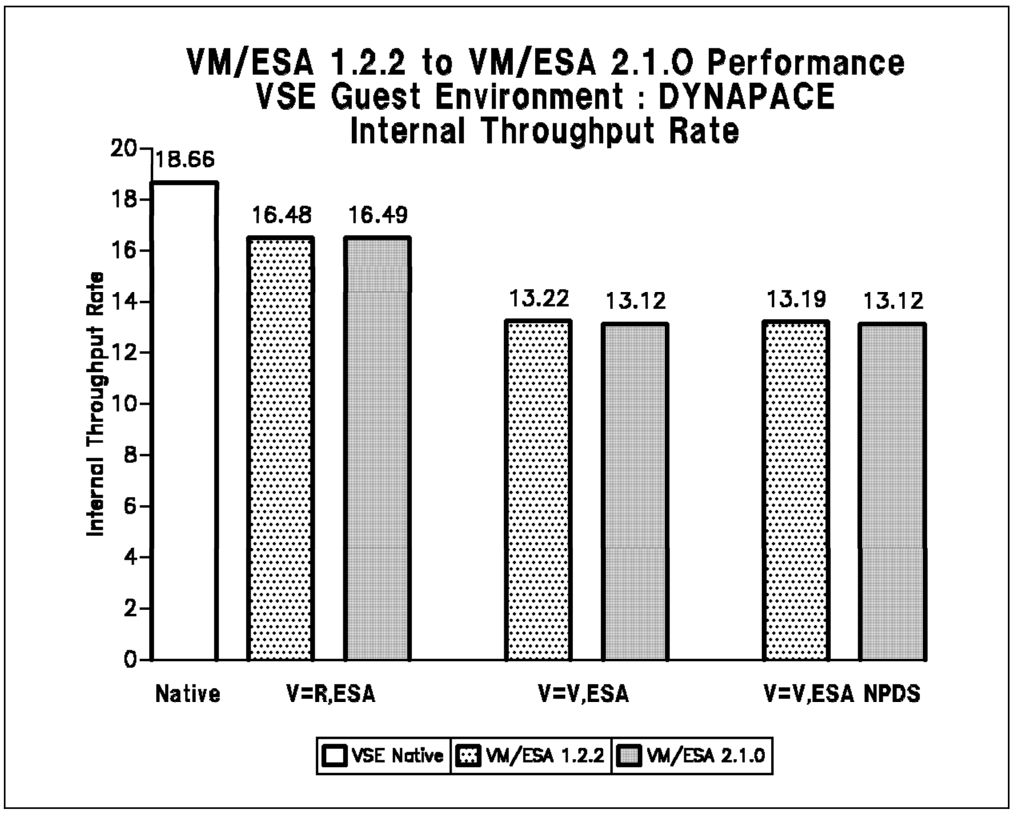

Figure 4. VSE guest migration from VM/ESA 1.2.2. DYNAPACE workload on a single VSE/ESA guest of VM/ESA 1.2.2 and VM/ESA 2.1.0 on the 9121-320 processor.

The VSE DYNAPACE workload was run as a guest of VM/ESA 1.2.2 and VM/ESA 2.1.0 in three modes: V=R, V=V, and V=V with the No Page Data Set (NPDS) option. The V=R guest environment had dedicated DASD with I/O assist. The two V=V guest environments were configured with full pack minidisk DASD with minidisk caching (MDC).

3 For more information on VSE/ESA 2.1.0 performance, refer to *VSE/ESA 2.1.0 Performance Considerations*.

For each guest environment, internal throughput rates were equivalent to VM/ESA 1.2.2 within measurement variability.

When comparing guest ITR to VSE native, the V=R guest, with this workload, achieves about 88% of native. The V=V guest achieves about 71% of native performance running with or without the NPDS option.

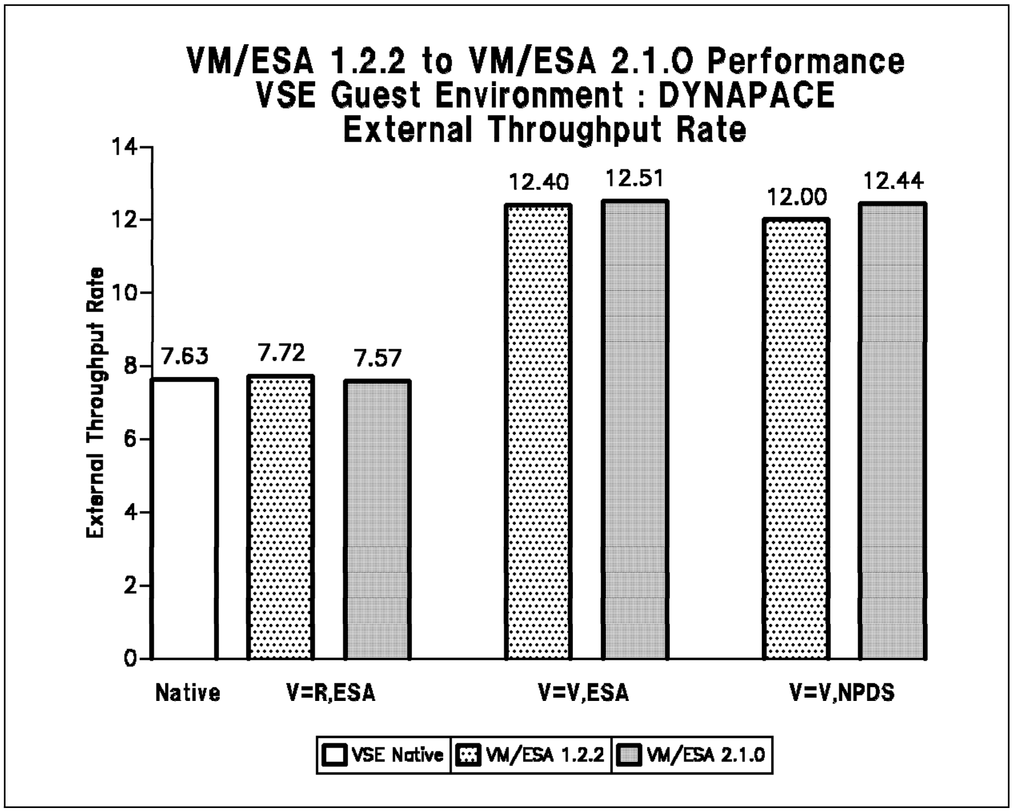

Figure 5. VSE guest migration from VM/ESA 1.2.2. DYNAPACE workload on a single VSE/ESA guest of VM/ESA 1.2.2 and VM/ESA 2.1.0 on the 9121-320 processor.

Figure 5 shows elapsed time comparisons of the various guest modes running under VM/ESA 1.2.2 and VM/ESA 2.1.0. The elapsed time duration of the batch jobs remains unchanged (within run variability) for all environments. The benefits of MDC are demonstrated in the ETR results for the V=V guest environments.

Workload: DYNAPACE

Hardware Configuration

Processor models: 9121-320[4]
Storage
Real: 256MB
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						1
3390-2	3990-02	4				10	2	
3380-K	3990-03	4				10		

Software Configuration

VSE version: 2.1.0 (using the Standard Dispatcher)

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
VSEVR	1	VSE V=R	96MB/ESA	100		IOASSIST ON CCWTRANS OFF
or						
VSEVV	1	VSE V=V	96MB/ESA	100		IOASSIST OFF
or						
VSEVV	1	VSE V=V NPDS	224MB/ESA	100		IOASSIST OFF
SMART	1	RTM	16MB/370	100		
WRITER	1	CP monitor	2MB/XA	100		

Additional Information: Starting with VM/ESA 1.2.2, minidisk caching (MDC) became available for use with non-CMS guests. Therefore, all V=V guest measurements in this section were run with MDC active.

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. The guest was run in three modes:

- V=R, mode=ESA
- V=V, mode=ESA
- V=V, mode=ESA NPDS (No Page Data-Set)

All DASD are dedicated to the VSE V=R guest for these measurements (except for the VM system DASD volumes). All V=V measurement environments were defined with full pack minidisks and MDC. The VM system used for these guest measurements has a 96MB V=R area defined. For measurements with V=V guests, the V=R area is configured, but not used. Therefore, if the real storage

[4] See “Hardware” on page 21 for an explanation of how this processor model was defined.

configuration on the processor is 256MB, then 160MB of 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 are provided in the following tables. The VSE native results are provided in Table 17 on page 69.

Table 8 (Page 1 of 2). VSE/ESA V=R guest migration from VM/ESA 1.2.2 on 9121-320, DYNAPACE workload.

VM/ESA Release Run ID	1.2.2 L1R78PF0	2.1.0 L1R88PF0	Difference	%Difference
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	870.0	888.0	18.0	2.07%
ETR (C)	7.72	7.57	-0.16	-2.03%
ITR (H)	16.48	16.49	0.01	0.04%
ITR	16.43	16.45	0.02	0.10%
ITRR (H)	1.000	1.000	0.000	0.04%
ITRR	1.000	1.001	0.001	0.10%
Proc. Usage (Sec)				
PBT/CMD (H)	3.640	3.639	-0.001	-0.04%
PBT/CMD	3.651	3.647	-0.004	-0.10%
CP/CMD (H)	0.275	0.281	0.007	2.50%
CP/CMD	0.233	0.238	0.005	2.07%
EMUL/CMD (H)	3.366	3.358	-0.008	-0.25%
EMUL/CMD	3.418	3.409	-0.009	-0.25%
Processor Util.				
TOTAL (H)	46.87	45.90	-0.97	-2.07%
TOTAL	47.00	46.00	-1.00	-2.13%
TOTAL EMUL (H)	43.33	42.35	-0.98	-2.27%
TOTAL EMUL	44.00	43.00	-1.00	-2.27%
TVR(H)	1.08	1.08	0.00	0.21%
TVR	1.07	1.07	0.00	0.15%
Storage				
NUCLEUS SIZE (V)	2584KB	2768KB	184KB	7.12%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	38624	38803	179	0.46%
FREEPGS	82	82	0	0.00%
FREE UTIL	0.55	0.55	0.00	0.55%
SHRPGS	1053	1052	-1	-0.09%
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

Table 8 (Page 2 of 2). VSE/ESA V=R guest migration from VM/ESA 1.2.2 on 9121-320, DYNAPACE workload.

VM/ESA Release Run ID	1.2.2 L1R78PF0	2.1.0 L1R88PF0	Difference	%Difference
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
I/O				
VIO RATE	1.000	1.000	0.000	0.00%
VIO/CMD	7.768	7.929	0.161	2.07%
RIO RATE (V)	2.000	2.000	0.000	0.00%
RIO/CMD (V)	15.536	15.857	0.321	2.07%
DASD IO TOTAL (V)	1583	1671	88	5.56%
DASD IO RATE (V)	1.76	1.86	0.10	5.56%
DASD IO/CMD (V)	13.66	14.72	1.06	7.74%
MDC REAL SIZE (MB)	7.3	6.0	-1.2	-17.10%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	0.03	0.03	0	0.00%
MDC WRITES (I/Os)	0.03	0.03	0	0.00%
MDC AVOID	0.01	0.01	0	0.00%
MDC HIT RATIO	0.40	0.30	-0.10	-25.00%
PRIVOPs				
PRIVOP/CMD (R)	10.884	10.884	0.000	0.00%
DIAG/CMD (R)	606.089	618.143	12.054	1.99%
SIE/CMD	2679.911	2671.929	-7.982	-0.30%
SIE INTCPT/CMD	2117.129	2110.824	-6.306	-0.30%
FREE TOTL/CMD	528.214	531.214	3.000	0.57%

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

Table 9 (Page 1 of 2). VSE/ESA V=V guest migration from VM/ESA 1.2.2 on 9121-320, DYNAPACE workload.

VM/ESA Release Run ID	1.2.2 L1V78PF0	2.1.0 L1V88PF0	Difference	%Difference
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	542.0	537.0	-5.0	-0.92%
ETR (C)	12.40	12.51	0.12	0.93%
ITR (H)	13.22	13.12	-0.09	-0.71%
ITR	13.93	13.17	-0.76	-5.44%
ITRR (H)	1.000	0.993	-0.007	-0.71%
ITRR	1.000	0.946	-0.054	-5.44%
Proc. Usage (Sec)				
PBT/CMD (H)	4.539	4.571	0.033	0.72%
PBT/CMD	4.307	4.555	0.248	5.76%
CP/CMD (H)	1.177	1.206	0.028	2.41%
CP/CMD	1.016	1.055	0.039	3.80%
EMUL/CMD (H)	3.362	3.366	0.004	0.12%
EMUL/CMD	3.291	3.500	0.209	6.36%
Processor Util.				
TOTAL (H)	93.79	95.35	1.55	1.66%
TOTAL	89.00	95.00	6.00	6.74%
TOTAL EMUL (H)	69.46	70.20	0.73	1.06%
TOTAL EMUL	68.00	73.00	5.00	7.35%
TVR(H)	1.35	1.36	0.01	0.59%
TVR	1.31	1.30	-0.01	-0.57%
Storage				
NUCLEUS SIZE (V)	2584KB	2768KB	184KB	7.12%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	38544	38719	175	0.45%
FREEPGS	101	104	3	2.97%
FREE UTIL	0.61	0.60	-0.01	-2.10%
SHRPGS	91	51	-40	-43.96%
Paging				
PAGE/CMD	82.268	177.402	95.134	115.64%
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	222.607	234.938	12.330	5.54%
I/O				
VIO RATE	610.000	651.000	41.000	6.72%
VIO/CMD	2951.964	3121.313	169.348	5.74%
RIO RATE (V)	270.000	272.000	2.000	0.74%
RIO/CMD (V)	1306.607	1304.143	-2.464	-0.19%
DASD IO TOTAL (V)	145053	146252	1199	0.83%
DASD IO RATE (V)	268.62	270.84	2.22	0.83%
DASD IO/CMD (V)	1299.91	1298.57	-1.35	-0.10%
MDC REAL SIZE (MB)	112.6	113.4	0.8	0.68%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	372	397	25	6.72%
MDC WRITES (I/Os)	185	198	13	7.03%
MDC AVOID	349	373	24	6.88%
MDC HIT RATIO	0.87	0.87	0	0.00%

Table 9 (Page 2 of 2). VSE/ESA V=V guest migration from VM/ESA 1.2.2 on 9121-320, DYNAPACE workload.

VM/ESA Release Run ID	1.2.2 L1V78PF0	2.1.0 L1V88PF0	Difference	%Difference
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD (R)	2949.616	3118.813	169.196	5.74%
DIAG/CMD (R)	458.144	471.902	13.757	3.00%
SIE/CMD	13027.357	13775.009	747.652	5.74%
SIE INTCPT/CMD	11333.801	11984.258	650.457	5.74%
FREE TOTL/CMD	3614.946	3787.768	172.821	4.78%

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

Table 10 (Page 1 of 2). VSE/ESA V=V NPDS guest migration from VM/ESA 1.2.2 on 9121-320 DYNAPACE workload.

VM/ESA Release Run ID	1.2.2 L1O78PF3	2.1.0 L1O88PF0	Difference	%Difference
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**224M**	**224M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	560.0	540.0	-20.0	-3.57%
ETR (C)	12.00	12.44	0.44	3.70%
ITR (H)	13.19	13.12	-0.07	-0.54%
ITR	13.19	13.10	-0.09	-0.66%
ITRR (H)	1.000	0.995	-0.005	-0.54%
ITRR	1.000	0.993	-0.007	-0.66%
Proc. Usage (Sec)				
PBT/CMD (H)	4.548	4.573	0.025	0.54%
PBT/CMD	4.550	4.580	0.030	0.67%
CP/CMD (H)	1.194	1.208	0.013	1.12%
CP/CMD	1.100	1.109	0.009	0.81%
EMUL/CMD (H)	3.354	3.365	0.011	0.33%
EMUL/CMD	3.450	3.471	0.021	0.62%

Table 10 (Page 2 of 2). VSE/ESA V=V NPDS guest migration from VM/ESA 1.2.2 on 9121-320 DYNAPACE workload.

VM/ESA Release Run ID	1.2.2 L1O78PF3	2.1.0 L1O88PF0	Difference	%Difference
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**224M**	**224M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Processor Util.				
TOTAL (H)	90.96	94.84	3.88	4.26%
TOTAL	91.00	95.00	4.00	4.40%
TOTAL EMUL (H)	67.08	69.79	2.72	4.05%
TOTAL EMUL	69.00	72.00	3.00	4.35%
TVR(H)	1.36	1.36	0.00	0.21%
TVR	1.32	1.32	0.00	0.05%
Storage				
NUCLEUS SIZE (V)	2584KB	2768KB	184KB	7.12%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	38767	38723	-44	-0.11%
FREEPGS	103	104	1	0.97%
FREE UTIL	0.61	0.60	-0.01	-2.15%
SHRPGS	317	143	-174	-54.89%
Paging				
PAGE/CMD	120.000	188.036	68.036	56.70%
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	235.000	236.250	1.250	0.53%
I/O				
VIO RATE	624.000	648.000	24.000	3.85%
VIO/CMD	3120.000	3124.286	4.286	0.14%
RIO RATE (V)	258.000	282.000	24.000	9.30%
RIO/CMD (V)	1290.000	1359.643	69.643	5.40%
DASD IO TOTAL (V)	153980	151742	-2238	-1.45%
DASD IO RATE (V)	256.63	281.00	24.37	9.50%
DASD IO/CMD (V)	1283.17	1354.84	71.67	5.59%
MDC REAL SIZE (MB)	112.9	113.2	0.2	0.19%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	381	395	14	3.67%
MDC WRITES (I/Os)	190	197	7	3.68%
MDC AVOID	358	372	14	3.91%
MDC HIT RATIO	0.87	0.87	0	0.00%
PRIVOPs				
PRIVOP/CMD (R)	3119.125	3120.384	1.259	0.04%
DIAG/CMD (R)	482.500	472.429	-10.071	-2.09%
SIE/CMD	13810.000	13760.357	-49.643	-0.36%
SIE INTCPT/CMD	12014.700	11971.511	-43.189	-0.36%
FREE TOTL/CMD	3795.000	3780.000	-15.000	-0.40%

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

9121-480 / DYNAPACE

This section examines VSE/ESA 2.1.0 guest performance of VM/ESA 1.2.2 compared to VM/ESA 2.1.0 running the DYNAPACE workload on a 9121-480. Because the 9121-480 has two processors, the VSE/ESA Turbo Dispatcher[5] is used.

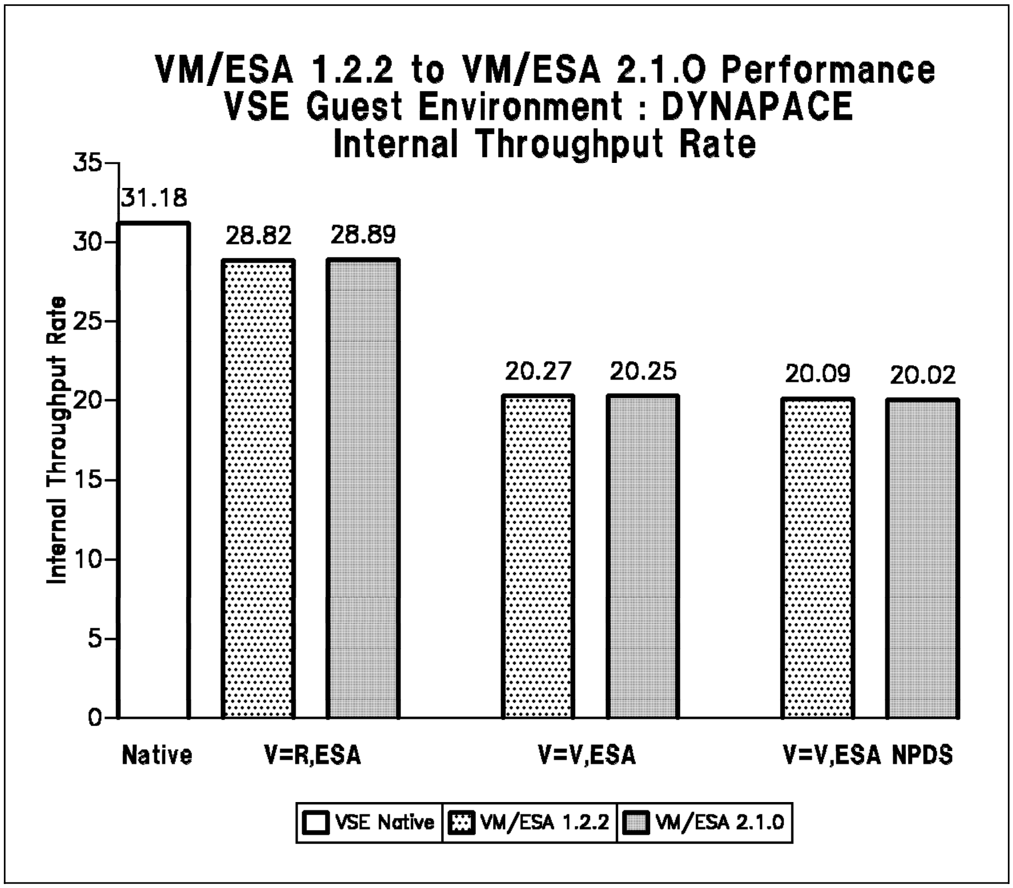

Figure 6. VSE guest migration from VM/ESA 1.2.2. DYNAPACE workload on a single VSE/ESA guest of VM/ESA 1.2.2 and VM/ESA 2.1.0 on the 9121-480 processor.

The VSE DYNAPACE workload was run as a guest of VM/ESA 1.2.2 and VM/ESA 2.1.0 in three modes: V=R, V=V, and V=V with the No Page Data Set (NPDS) option. The V=R guest environment had dedicated DASD with I/O assist. The two V=V guest environments were configured with full pack minidisk DASD with minidisk caching (MDC). The VSE/ESA Turbo Dispatcher was enabled with 2 processors active. All Turbo environments were run with un-dedicated processors.

For each guest mode, internal throughput rates were equivalent to VM/ESA 1.2.2 within measurement variability.

5 For more information on the VSE/ESA Turbo Dispatcher, refer to *VSE/ESA 2.1 Turbo Dispatcher Performance*.

When comparing guest ITR to VSE native, the V=R guest, with this workload, achieves about 91% of native. The V=V guest achieves about 63% of native performance running with or without the NPDS option.

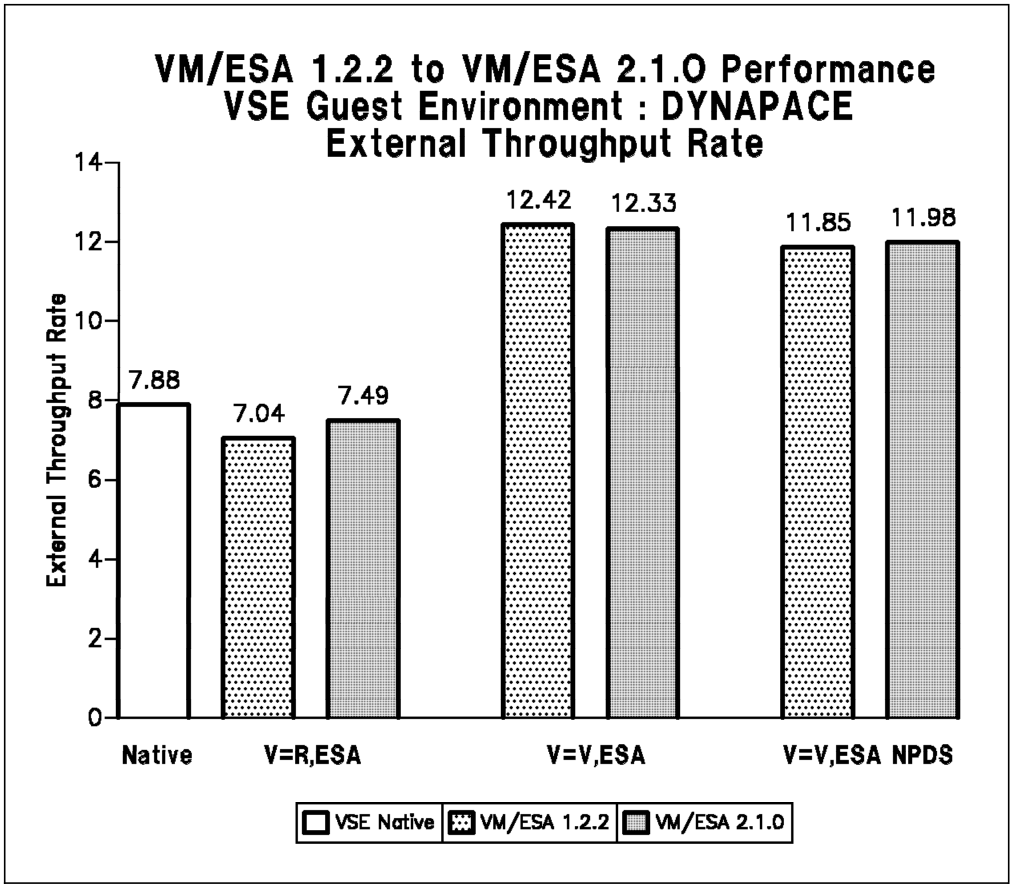

Figure 7. VSE guest migration from VM/ESA 1.2.2. DYNAPACE workload on a single VSE/ESA guest of VM/ESA 1.2.2 and VM/ESA 2.1.0 on the 9121-480 processor.

Figure 7 shows elapsed time comparisons of the various guest modes running under VM/ESA 1.2.2 and VM/ESA 2.1.0. The elapsed time duration of the batch jobs remains unchanged (within run variability) for all environments. The benefits of MDC are demonstrated in the ETR results for the V=V guest environments.

Workload: DYNAPACE

Hardware Configuration

Processor models: 9121-480
Storage
Real: 256MB
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						1
3390-2	3990-02	4				10	2	
3380-K	3990-03	4				10		

Software Configuration

VSE version: 2.1.0 (using the Turbo Dispatcher)

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
VSEVR	1	VSE V=R	96MB/ESA	100		IOASSIST ON CCWTRANS OFF UNDEDICATE
or						
VSEVV	1	VSE V=V	96MB/ESA	100		IOASSIST OFF
or						
VSEVV	1	VSE V=V NPDS	224MB/ESA	100		IOASSIST OFF
SMART	1	RTM	16MB/370	100		
WRITER	1	CP monitor	2MB/XA	100		

Additional Information: Starting with VM/ESA 1.2.2, minidisk caching (MDC) became available for use with non-CMS guests. Therefore, all V=V guest measurements in this section were run with MDC active.

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. The guest was run in three modes:

- V=R, mode=ESA
- V=V, mode=ESA
- V=V, mode=ESA NPDS (No Page Data-Set)

All DASD are dedicated to the VSE V=R guest for these measurements (except for the VM system DASD volumes). All V=V measurement environments were defined with full pack minidisks and MDC. The VM system used for these guest measurements has a 96MB V=R area defined. For measurements with V=V guests, the V=R area is configured, but not used. Therefore, if the real storage configuration on the processor is 256MB, then 160MB of 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 are provided in the following tables. The VSE native results are provided in Table 17 on page 69.

Table 11 (Page 1 of 2). VSE/ESA V=R guest migration from VM/ESA 1.2.2 on 9121-480, DYNAPACE workload.

VM/ESA Release **Run ID**	**1.2.2** **L2R78PF2**	**2.1.0** **L2R88PF1**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**2**	**2**		
Throughput (Min)				
Elapsed Time (C)	954.0	897.0	-57.0	-5.97%
ETR (C)	7.04	7.49	0.45	6.35%
ITR (H)	28.82	28.89	0.07	0.24%
ITR	28.75	28.81	0.06	0.22%
ITRR (H)	1.000	1.002	0.002	0.24%
ITRR	1.000	1.002	0.002	0.22%
Proc. Usage (Sec)				
PBT/CMD (H)	4.163	4.154	-0.010	-0.23%
PBT/CMD	4.174	4.165	-0.009	-0.22%
CP/CMD (H)	0.401	0.398	-0.003	-0.82%
CP/CMD	0.341	0.400	0.060	17.53%
EMUL/CMD (H)	3.762	3.756	-0.006	-0.17%
EMUL/CMD	3.833	3.764	-0.069	-1.80%
Processor Util.				
TOTAL (H)	48.88	51.86	2.98	6.10%
TOTAL	49.00	52.00	3.00	6.12%
UTIL/PROC (H)	24.44	25.93	1.49	6.10%
UTIL/PROC	24.50	26.00	1.50	6.12%
TOTAL EMUL (H)	44.17	46.89	2.73	6.17%
TOTAL EMUL	45.00	47.00	2.00	4.44%
MASTER TOTAL (H)	21.42	23.93	2.51	11.70%
MASTER TOTAL	21.00	24.00	3.00	14.29%
MASTER EMUL (H)	19.71	22.06	2.34	11.89%
MASTER EMUL	20.00	22.00	2.00	10.00%
TVR(H)	1.11	1.11	0.00	-0.06%
TVR	1.09	1.11	0.02	1.61%
Storage				
NUCLEUS SIZE (V)	2584KB	2776KB	192KB	7.43%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
PGBLPGS	38559	38528	-31	-0.08%
FREEPGS	89	87	-2	-2.25%
FREE UTIL	0.54	0.55	0.01	1.24%
SHRPGS	1054	1054	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

Table 11 (Page 2 of 2). VSE/ESA V=R guest migration from VM/ESA 1.2.2 on 9121-480, DYNAPACE workload.

VM/ESA Release Run ID	1.2.2 L2R78PF2	2.1.0 L2R88PF1	Difference	%Difference
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**2**	**2**		
I/O				
VIO RATE	1.000	1.000	0.000	0.00%
VIO/CMD	8.518	8.009	-0.509	-5.97%
RIO RATE (V)	2.000	2.000	0.000	0.00%
RIO/CMD (V)	17.036	16.018	-1.018	-5.97%
DASD IO TOTAL (V)	1744	1532	-212	-12.16%
DASD IO RATE (V)	1.82	1.82	0.01	0.39%
DASD IO/CMD (V)	15.47	14.61	-0.87	-5.61%
MDC REAL SIZE (MB)	6.0	5.9	-0.1	-1.94%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	0.03	0.03	0	0.00%
MDC WRITES (I/Os)	0.01	0.01	0	0.00%
MDC AVOID	0.00	0.00	0	na
MDC HIT RATIO	0.27	0.26	-0.01	-3.70%
PRIVOPs				
PRIVOP/CMD (R)	11.435	11.009	-0.426	-3.72%
DIAG/CMD (R)	668.750	641.482	-27.268	-4.08%
SIE/CMD	2785.339	2715.027	-70.313	-2.52%
SIE INTCPT/CMD	2172.565	2090.571	-81.994	-3.77%
FREE TOTL/CMD	596.250	544.607	-51.643	-8.66%

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

Table 12 (Page 1 of 2). VSE/ESA V=V guest migration from VM/ESA 1.2.2 on 9121-480, DYNAPACE workload.

VM/ESA Release **Run ID**	**1.2.2** **L2V78PF2**	**2.1.0** **L2V88PF3**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**2**	**2**		
Throughput (Min)				
Elapsed Time (C)	541.0	545.0	4.0	0.74%
ETR (C)	12.42	12.33	-0.09	-0.73%
ITR (H)	20.27	20.25	-0.01	-0.05%
ITR	20.36	20.21	-0.15	-0.73%
ITRR (H)	1.000	0.999	-0.001	-0.05%
ITRR	1.000	0.993	-0.007	-0.73%
Proc. Usage (Sec)				
PBT/CMD (H)	5.922	5.925	0.003	0.05%
PBT/CMD	5.893	5.937	0.044	0.74%
CP/CMD (H)	1.665	1.688	0.023	1.38%
CP/CMD	1.497	1.557	0.060	3.99%
EMUL/CMD (H)	4.257	4.236	-0.020	-0.47%
EMUL/CMD	4.396	4.379	-0.016	-0.37%
Processor Util.				
TOTAL (H)	122.59	121.75	-0.84	-0.68%
TOTAL	122.00	122.00	0.00	0.00%
UTIL/PROC (H)	61.30	60.88	-0.42	-0.68%
UTIL/PROC	61.00	61.00	0.00	0.00%
TOTAL EMUL (H)	88.12	87.06	-1.06	-1.20%
TOTAL EMUL	91.00	90.00	-1.00	-1.10%
MASTER TOTAL (H)	61.55	61.10	-0.45	-0.72%
MASTER TOTAL	61.00	61.00	0.00	0.00%
MASTER EMUL (H)	44.67	44.18	-0.50	-1.11%
MASTER EMUL	46.00	46.00	0.00	0.00%
TVR(H)	1.39	1.40	0.01	0.52%
TVR	1.34	1.36	0.01	1.11%
Storage				
NUCLEUS SIZE (V)	2584KB	2768KB	184KB	7.12%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
PGBLPGS	38487	38456	-31	-0.08%
FREEPGS	107	108	1	0.93%
FREE UTIL	0.60	0.58	-0.02	-3.62%
SHRPGS	201	331	130	64.68%
Paging				
PAGE/CMD	115.929	92.455	-23.473	-20.25%
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	318.804	321.161	2.357	0.74%

Table 12 (Page 2 of 2). VSE/ESA V=V guest migration from VM/ESA 1.2.2 on 9121-480, DYNAPACE workload.

VM/ESA Release **Run ID**	**1.2.2** **L2V78PF2**	**2.1.0** **L2V88PF3**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**2**	**2**		
I/O				
VIO RATE	645.000	642.000	-3.000	-0.47%
VIO/CMD	3115.580	3124.018	8.438	0.27%
RIO RATE (V)	281.000	275.000	-6.000	-2.14%
RIO/CMD (V)	1357.330	1338.170	-19.161	-1.41%
DASD IO TOTAL (V)	151432	148149	-3283	-2.17%
DASD IO RATE (V)	280.43	274.30	-6.13	-2.18%
DASD IO/CMD (V)	1354.58	1322.76	-31.82	-2.35%
MDC REAL SIZE (MB)	112.5	111.6	-0.9	-0.77%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	395	393	-2	-0.51%
MDC WRITES (I/Os)	196	195	-1	-0.51%
MDC AVOID	369	371	2	0.54%
MDC HIT RATIO	0.86	0.87	0.01	1.16%
PRIVOPs				
PRIVOP/CMD (R)	3120.286	3120.393	0.107	0.00%
DIAG/CMD (R)	728.357	715.125	-13.232	-1.82%
SIE/CMD	16457.027	16637.098	180.071	1.09%
SIE INTCPT/CMD	15140.465	15139.759	-0.705	0.00%
FREE TOTL/CMD	3791.830	3756.607	-35.223	-0.93%

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

Table 13 (Page 1 of 2). VSE/ESA V=V NPDS guest migration from VM/ESA 1.2.2 on 9121-480 DYNAPACE workload.				
VM/ESA Release **Run ID**	**1.2.2** **L2O78PF2**	**2.1.0** **L2O88PF2**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**224M**	**224M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**2**	**2**		
Throughput (Min)				
Elapsed Time (H)	567.0	561.0	-6.0	-1.06%
ETR (H)	11.85	11.98	0.13	1.07%
ITR (H)	20.09	20.02	-0.08	-0.38%
ITR	20.09	19.96	-0.12	-0.61%
ITRR (H)	1.000	0.996	-0.004	-0.38%
ITRR	1.000	0.994	-0.006	-0.61%
Proc. Usage (Sec)				
PBT/CMD (H)	5.972	5.995	0.023	0.39%
PBT/CMD	5.974	6.011	0.037	0.62%
CP/CMD (H)	1.684	1.684	0.000	-0.02%
CP/CMD	1.519	1.553	0.034	2.24%
EMUL/CMD (H)	4.288	4.311	0.023	0.54%
EMUL/CMD	4.455	4.458	0.003	0.07%
Processor Util.				
TOTAL (H)	117.96	119.69	1.72	1.46%
TOTAL	118.00	120.00	2.00	1.69%
UTIL/PROC (H)	58.98	59.84	0.86	1.46%
UTIL/PROC	59.00	60.00	1.00	1.69%
TOTAL EMUL (H)	84.70	86.07	1.37	1.62%
TOTAL EMUL	88.00	89.00	1.00	1.14%
MASTER TOTAL (H)	58.43	60.04	1.61	2.76%
MASTER TOTAL	58.00	60.00	2.00	3.45%
MASTER EMUL (H)	42.47	43.70	1.24	2.91%
MASTER EMUL	44.00	45.00	1.00	2.27%
TVR(H)	1.39	1.39	0.00	-0.16%
TVR	1.34	1.35	0.01	0.55%
Storage				
NUCLEUS SIZE (V)	2584KB	2768KB	184KB	7.12%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
PGBLPGS	38484	38455	-29	-0.08%
FREEPGS	107	109	2	1.87%
FREE UTIL	0.62	0.59	-0.03	-4.81%
SHRPGS	37	363	326	881.08%
Paging				
PAGE/CMD	207.563	105.187	-102.375	-49.32%
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	243.000	315.562	72.562	29.86%

Table 13 (Page 2 of 2). VSE/ESA V=V NPDS guest migration from VM/ESA 1.2.2 on 9121-480 DYNAPACE workload.

VM/ESA Release **Run ID**	**1.2.2** **L2O78PF2**	**2.1.0** **L2O88PF2**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**224M**	**224M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**2**	**2**		
I/O				
VIO RATE	617.000	622.000	5.000	0.81%
VIO/CMD	3123.563	3115.554	-8.009	-0.26%
RIO RATE (V)	270.000	275.000	5.000	1.85%
RIO/CMD (V)	1366.875	1377.455	10.580	0.77%
DASD IO TOTAL (V)	145094	147802	2708	1.87%
DASD IO RATE (V)	268.69	273.71	5.01	1.87%
DASD IO/CMD (V)	1360.26	1370.98	10.72	0.79%
MDC REAL SIZE (MB)	112.1	112.3	0.1	0.13%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	377	381	4	1.06%
MDC WRITES (I/Os)	187	189	2	1.07%
MDC AVOID	355	359	4	1.13%
MDC HIT RATIO	0.87	0.87	0	0.00%
PRIVOPs				
PRIVOP/CMD (R)	3118.402	3119.393	0.991	0.03%
DIAG/CMD (R)	768.429	756.187	-12.241	-1.59%
SIE/CMD	16478.438	16489.393	10.955	0.07%
SIE INTCPT/CMD	15160.163	15170.241	10.079	0.07%
FREE TOTL/CMD	3766.500	3771.723	5.223	0.14%

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

9121-320 / VSECICS

This section examines VSE/ESA 2.1.0 guest performance running under VM/ESA 2.1.0 compared to VM/ESA 1.2.2. 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 157 for a detailed description of the workload. All DASD are dedicated to the VSE V=R guest for these measurements (except for the VM system DASD volumes). All V=V guest measurements use full pack minidisk DASD and minidisk caching.

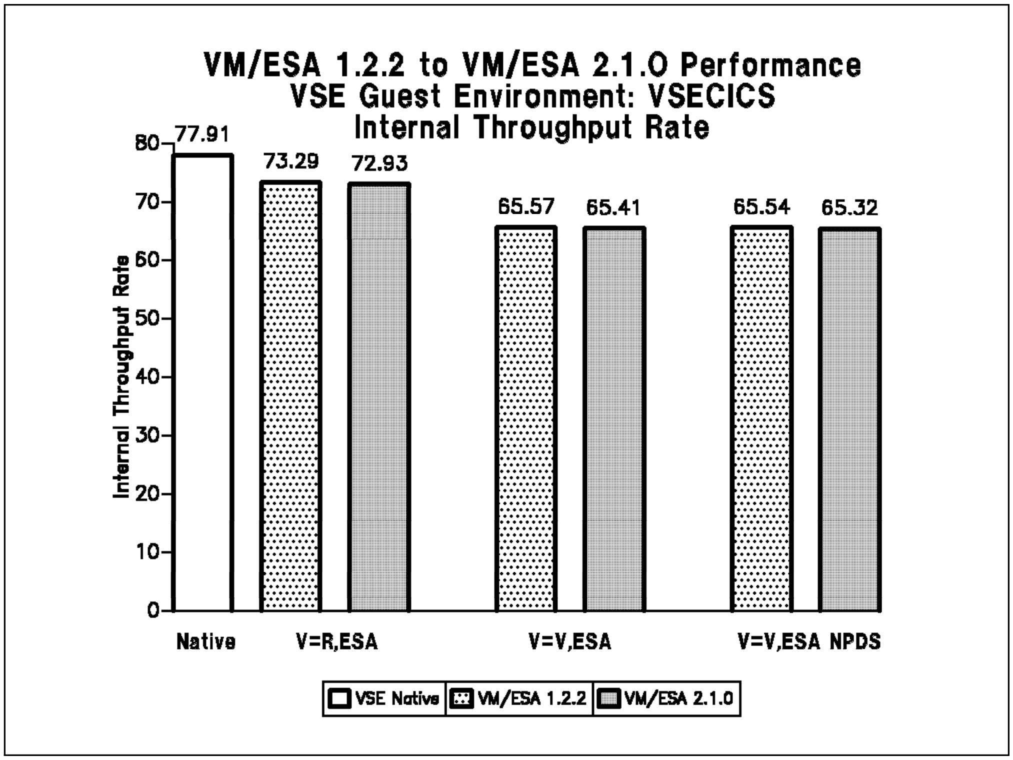

Figure 8. VSE guest migration from VM/ESA 1.2.2. VSECICS workload on a single VSE/ESA guest of VM/ESA 1.2.2 and VM/ESA 2.1.0 on the 9121-320 processor.

Comparing VM/ESA 1.2.2 to VM/ESA 2.1.0, internal throughput rates were equivalent within measurement variability.

When comparing guest ITR to VSE running native, the V=R guest, with this workload, achieved about 94% of native. The V=V guest achieved approximately 84% of native mode performance running with or without the NPDS option.

Workload: VSECICS

Hardware Configuration

Processor models: 9121-320[6]
Storage
Real: 256MB
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						2
3380-K	3990-03	4				16		
3390-2	3990-02	4				20	2	

Software Configuration

VSE version: 2.1.0 (using the Standard Dispatcher)

Virtual Machines:

Virtual Machine	Number	Type	Machine Size/Mode	SHARE	RESERVED	Other Options
VSEVR	1	VSE V=R	96MB/ESA	100		IOASSIST ON CCWTRANS OFF
or						
CICVV	1	VSE V=V	96MB/ESA	100		IOASSIST OFF
or						
CICVV	1	VSE V=V NPDS	300MB/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 1.2.2 and VM/ESA 2.1.0 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 1.2.2, the CPU utilization was near 90%. Then, the same number of terminals was run in the same guest mode under VM/ESA 2.1.0.

The VSE guest measurement results for the VSECICS workload are provided in the following tables. The VSE native results are provided in Table 17 on page 69.

[6] See "Hardware" on page 21 for an explanation of how this processor model was defined.

Table 14. VSE/ESA V=R guest migration from VM/ESA 1.2.2 on 9121-320, VSECICS.

VM/ESA Release Run ID	1.2.2 L1R78C90	2.1.0 L1R88C90	Difference	%Difference
Environment				
Real Storage	**256M**	**256M**		
Exp. Storage	**0M**	**0M**		
Users	**0960**	**0960**		
VM Size	**96M**	**96M**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Response Time				
AVG RESP (C)	0.226	0.226	0.000	-0.10%
Throughput				
ETR (C)	65.79	65.76	-0.02	-0.03%
ITR (H)	73.29	72.93	-0.36	-0.49%
ITRR (H)	1.000	0.995	-0.005	-0.49%
Proc. Usage				
PBT/CMD (H)	13.644	13.712	0.068	0.50%
PBT/CMD	13.681	13.685	0.005	0.03%
CP/CMD (H)	0.416	0.429	0.013	3.10%
CP/CMD	0.456	0.304	-0.152	-33.31%
EMUL/CMD (H)	13.228	13.283	0.055	0.41%
EMUL/CMD	13.225	13.381	0.157	1.18%
Processor Util.				
TOTAL (H)	89.76	90.17	0.42	0.46%
TOTAL	90.00	90.00	0.00	0.00%
TOTAL EMUL (H)	87.02	87.35	0.33	0.38%
TOTAL EMUL	87.00	88.00	1.00	1.15%
TVR(H)	1.03	1.03	0.00	0.08%
TVR	1.03	1.02	-0.01	-1.14%
Storage				
NUCLEUS SIZE (V)	2584K	2768K	184K	7.12%
TRACE TABLE (V)	200K	200K	0K	0.00%
PGBLPGS	38626	38800	174	0.45%
FREEPGS	80	86	6	7.50%
FREE UTIL	0.57	0.53	-0.04	-6.47%
SHRPGS	1080	1052	-28	-2.59%
Paging				
PAGE/CMD	0.000	0.000	0.000	na
XSTOR/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)	2.000	2.000	0.000	na
RIO/CMD (V)	0.030	0.030	0.000	na
MDC REAL SIZE (MB)	7.7	8.2	0.5	6.10%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	0.03	0.03	0	0.00%
MDC WRITES (I/Os)	0.03	0.03	0	0.00%
MDC AVOID	0.01	0.01	0	0.00%
MDC HIT RATIO	0.40	0.30	-0.10	-25.00%
PRIVOPs				
PRIVOP/CMD (R)	0.008	0.008	0.000	-1.26%
DIAG/CMD (R)	0.694	0.694	0.000	0.03%
SIE/CMD	4.423	4.455	0.032	0.72%
SIE INTCPT/CMD	2.256	2.272	0.016	0.72%
FREE TOTL/CMD	0.882	0.897	0.016	1.76%

Note: V=VMPRF, H=Hardware Monitor, C=CICSPARS, Unmarked=RTM

Table 15. VSE/ESA V=V guest migration from VM/ESA 1.2.2 on 9121-320, VSECICS.

VM/ESA Release **Run ID**	**1.2.2** **L1V78C90**	**2.1.0** **L1V88C90**	**Difference**	**%Difference**
Environment				
Real Storage	**160M**	**160M**		
Exp. Storage	**0M**	**0M**		
Users	**0880**	**0880**		
VM Size	**96M**	**96M**		
Guset Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Response Time				
AVG RESP (C)	0.239	0.237	-0.002	-0.79%
Throughput				
ETR (C)	60.20	60.21	0.01	0.02%
ITR (H)	65.58	65.42	-0.16	-0.25%
ITRR (H)	1.000	0.998	-0.002	-0.25%
Proc. Usage				
PBT/CMD (H)	15.249	15.287	0.037	0.25%
PBT/CMD	15.283	15.281	-0.003	-0.02%
CP/CMD (H)	1.799	1.828	0.029	1.64%
CP/CMD	1.661	1.661	0.000	-0.02%
EMUL/CMD (H)	13.451	13.459	0.008	0.06%
EMUL/CMD	13.622	13.620	-0.002	-0.02%
Processor Util.				
TOTAL (H)	91.80	92.04	0.24	0.26%
TOTAL	92.00	92.00	0.00	0.00%
TOTAL EMUL (H)	80.97	81.03	0.06	0.08%
TOTAL EMUL	82.00	82.00	0.00	0.00%
TVR(H)	1.13	1.14	0.00	0.19%
TVR	1.12	1.12	0.00	0.00%
Storage				
NUCLEUS SIZE (V)	2584K	2768K	184K	7.12%
TRACE TABLE (V)	200K	200K	0K	0.00%
PGBLPGS	38768	38718	-50	-0.13%
FREEPGS	103	107	4	3.88%
FREE UTIL	0.67	0.64	-0.03	-5.13%
SHRPGS	36	36	0	0.00%
Paging				
PAGE/CMD	0.017	0.000	-0.017	-100.00%
XSTOR/CMD	0.000	0.000	0.000	na
I/O				
VIO RATE	200.000	199.000	-1.000	-0.50%
VIO/CMD	3.322	3.305	-0.017	-0.52%
RIO RATE (V)	130.000	127.000	-3.000	-2.31%
RIO/CMD (V)	2.160	2.109	-0.050	-2.32%
MDC REAL SIZE (MB)	113.1	113.2	0.1	0.13%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	86	83	-3	-3.49%
MDC WRITES (I/Os)	67	70	3	4.48%
MDC AVOID	70	73	3	4.29%
MDC HIT RATIO	0.76	0.82	0.06	7.89%
PRIVOPs				
PRIVOP/CMD (R)	3.326	3.316	-0.010	-0.29%
DIAG/CMD (R)	0.782	0.782	-0.001	-0.09%
SIE/CMD	18.174	18.187	0.014	0.07%
SIE INTCPT/CMD	14.176	14.186	0.011	0.07%
FREE TOTL/CMD	4.950	4.684	-0.267	-5.38%

Note: V=VMPRF, H=Hardware Monitor, C=CICSPARS, Unmarked=RTM

Table 16. VSE/ESA V=V NPDS guest migration from VM/ESA 1.2.2 on 9121-320, VSECICS.

VM/ESA Release **Run ID**	**1.2.2** **L1O78C90**	**2.1.0** **L1O88C90**	**Difference**	**%Difference**
Environment				
Real Storage	**160M**	**160M**		
Exp. Storage	**0M**	**0M**		
Users	**0880**	**0880**		
VM Size	**300M**	**300M**		
Guset Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Response Time				
AVG RESP (C)	0.240	0.241	0.001	0.37%
Throughput				
ETR (C)	60.07	60.10	0.03	0.05%
ITR (H)	65.54	65.32	-0.22	-0.33%
ITRR (H)	1.000	0.997	-0.003	-0.33%
Proc. Usage				
PBT/CMD (H)	15.258	15.309	0.051	0.33%
PBT/CMD	15.316	15.309	-0.007	-0.05%
CP/CMD (H)	1.805	1.820	0.014	0.80%
CP/CMD	1.665	1.664	-0.001	-0.05%
EMUL/CMD (H)	13.452	13.489	0.036	0.27%
EMUL/CMD	13.651	13.645	-0.006	-0.05%
Processor Util.				
TOTAL (H)	91.65	92.00	0.35	0.38%
TOTAL	92.00	92.00	0.00	0.00%
TOTAL EMUL (H)	80.81	81.06	0.26	0.32%
TOTAL EMUL	82.00	82.00	0.00	0.00%
TVR(H)	1.13	1.13	0.00	0.06%
TVR	1.12	1.12	0.00	0.00%
Storage				
NUCLEUS SIZE (V)	2584K	2768K	184K	7.12%
TRACE TABLE (V)	200K	200K	0K	0.00%
PGBLPGS	38763	38724	-39	-0.10%
FREEPGS	104	103	-1	-0.96%
FREE UTIL	0.66	0.66	-0.01	-0.86%
SHRPGS	36	36	0	0.00%
Paging				
PAGE/CMD	0.017	0.017	0.000	-0.05%
XSTOR/CMD	0.000	0.000	0.000	na
I/O				
VIO RATE	199.000	199.000	0.000	0.00%
VIO/CMD	3.313	3.311	-0.002	-0.05%
RIO RATE (V)	130.000	127.000	-3.000	-2.31%
RIO/CMD (V)	2.164	2.113	-0.051	-2.35%
MDC REAL SIZE (MB)	113.0	113.1	0.1	0.12%
MDC READS (I/Os)	87	83	-4	-4.60%
MDC WRITES (I/Os)	67	69	2	2.99%
MDC AVOID	70	72	2	2.86%
MDC HIT RATIO	0.75	0.81	0.06	8.00%
PRIVOPs				
PRIVOP/CMD (R)	3.323	3.323	0.000	0.01%
DIAG/CMD (R)	0.783	0.784	0.001	0.06%
SIE/CMD	18.179	18.171	-0.008	-0.05%
SIE INTCPT/CMD	14.180	14.173	-0.007	-0.05%
FREE TOTL/CMD	4.978	4.709	-0.269	-5.39%

Note: V=VMPRF, H=Hardware Monitor, C=CICSPARS, Unmarked=RTM

Table 17. VSE/ESA 2.1.0 native results.

VSE/ESA Release	**2.1.0**	**2.1.0**	**2.1.0**
Workload	**DYNAPACE**	**DYNAPACE**	**VSECICS**
Processor	**9121-320**	**9121-480**	**9121-320**
Run ID	**L1N_8PF0**	**L2N_8PF0**	**L1N_8C90**
Environment			
Real Storage	**96M**	**96M**	**96M**
Users	**na**	**na**	**1040**
Processors	**1**	**2**	**1**
Response Time			
AVG RESP (C)	na	na	0.230
Throughput			
ETR (C)	7.63	7.88	71.22
ITR (H)	18.66	31.18	77.91
Proc. Usage			
PBT/CMD (H)	3.216	3.847	0.0128
Processor Util.			
TOTAL (H)	40.88	50.52	91.41
UTIL/PROC (H)	40.88	25.26	91.41

Note: H=Hardware Monitor, C=VSE console (DYNAPACE), C=CICSPARS (VSECICS)

VMSES/E

VM Service Enhancements Staged/Extended (VMSES/E) in VM/ESA 2.1.0 includes a number of performance enhancements. Some of these improved execution performance:

- VMFBLD was improved in a number of areas (such as the new CSLGEN option of NUC, the improved performance and storage handling of requisite processing, and the minimization of multiple part handler calls, while still respecting requisite order).
- By adding the SPRODID option to VMFCOPY, the selection of files to be copied can be controlled.

Other VMSES/E 2.1.0 enhancements reduced the number of manual steps via automation:

- The new CNTRL option to override control file name in the PPF (for GENCPBLS, VMFEXUPD, VMFNLS, VMFxASM family of commands).
- The build list option CNTRL has been added to VMFBDGEN to override CNTRL in the PPF.
- To allow multiple products the ability to update the CP load list, the GENCPBLS command has been modified to process multiple xxxMDLAT MACROs.

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 2.1.0 came from the VMFBLD improvements. The regression measurements showed significant savings in build response time.

The improvements in the build function result in a virtual storage reduction by storing global dependencies only once. By removing duplicate object processing, dependency and requisite processing is bypassed.[7] The VMFE2E module has been restructured and optimized to perform both chained GETs and SETs for data and to use its buffers more efficiently. VMFSIMPC has been modified to directly return data back two levels rather than use VMFSIM as an intermediate step when returning stem data back to VMFBLD (and other functions). Finally, VMFMSG has been improved by grouping its VMFE2E calls into one invocation (which helps VMFBLD).

[7] This portion of the VMFBLD improvements is also available on VM/ESA 1.2.2 as APAR VM57938.

Overall, for the dedicated, single-user measurements reported here, the process of receiving and applying CMS service, and building CMS on VMSES/E 2.1.0 showed total elapsed time improved 6% when compared to VMSES/E 1.2.2. Virtual CPU time and Total CPU time each improved 13%. These measurements were made on the 9121-480 configuration.

The following measurements are provided to demonstrate the performance of these changes on VMSES/E 2.1.0.

Hardware Configuration

Processor model:	9121-480
Processors used:	2
Storage	
Real:	256MB (default MDC)
Expanded:	0MB
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*
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.

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 2.1.0 did not degrade performance when compared to VMSES/E 1.2.2 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 1.2.2 and VMSES/E 2.1.0; 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 2.1.0 and VMSES/E 1.2.2. 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.)

Workload: Receive

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

Table 18. VMFREC measurement data: migration from VMSES/E 1.2.2 on the 9121-480

VMSES/E Release	1.2.2	2.1.0	Difference	%Difference
Total Time (QT)	557	553	-4	-1%
Total CPU (QT)	187	182	-5	-2%
Virtual CPU (QT)	171	167	-4	-2%

Note: QT=CP QUERY TIME

Workload: Apply

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

Table 19. VMFAPPLY measurement data: migration from VMSES/E 1.2.2 on the 9121-480

VMSES/E Release	1.2.2	2.1.0	Difference	%Difference
Total Time (QT)	359	349	-10	-3%
Total CPU (QT)	286	277	-9	-3%
Virtual CPU (QT)	279	271	-8	-3%

Note: QT=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: migration from VMSES/E 1.2.2 on the 9121-480

VMSES/E Release	1.2.2	2.1.0	Difference	%Difference
Total Time (QT)	135	103	-32	-24%
Total CPU (QT)	119	88	-31	-26%
Virtual CPU (QT)	118	87	-31	-26%

Note: QT=CP QUERY TIME

Workload: Build with SERVICED Option

Command: VMFBLD PPF ESA CMS (SERVICED
Scenario Details: 15 build lists processed; 149 objects built.

Table 21. VMFBLD SERVICED measurement data: migration from VMSES/E 1.2.2 on the 9121-480

VMSES/E Release	1.2.2	2.1.0	Difference	%Difference
Total Time (QT)	667	606	-61	-10%
Total CPU (QT)	311	239	-72	-23%
Virtual CPU (QT)	298	227	-71	-24%

Note: QT=CP QUERY TIME

Migration from Other VM Releases

The performance results provided in this report apply to migration from VM/ESA 1.2.2. 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.

Note: In this section, VM/ESA releases prior to VM/ESA 2.1.0 are sometimes referred to without the version number. For example, VM/ESA 2.2 refers to VM/ESA Version 1 Release 2.2.

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*
- **210** *VM/ESA Version 2 Release 1.0 Performance Report* (this document)

See "Referenced Publications" on page 5 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 1.2.1 to VM/ESA 2.1.0 will tend to affect VSE guest performance, look at the VM/ESA 1.2.1 to VM/ESA 1.2.2 comparison measurements and the VM/ESA 1.2.2 to VM/ESA 2.1.0 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 79.

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 76.

Table 22. 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			
VM/ESA 2.2	VM/ESA 2.1.0	9121-742	210			
		9121-480	210		210	
		9121-320			210	
		9221-170	210			

Migration Summary: CMS-Intensive Environment

A large body of performance information for the CMS-intensive environment has been 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 74.

Table 23 summarizes the ITR relationships that were observed for the CMS-intensive environment for a number of VM release-to-release transitions:

Table 23. Approximate VM relative capacity: CMS-intensive environment

Source	Target	Case	ITRR	ITRR Derivation	Notes
VM/SP 5	VM/ESA 1.5 (370) VM/ESA 2.1.0	 9221-120	0.94 0.90	R5*R13c R5*R13a*R2*R21*R22	1,5,7 1,2,4,6-8
VM/SP 6	VM/ESA 1.5 (370) VM/ESA 2.1.0	 9221-120	1.09 1.05	R6*R13c R6*R13a*R2*R21*R22	5 2,4,6-8
VM/ESA 1.0 (370)	VM/ESA 1.5 (370) VM/ESA 2.1.0	 9221-120 9221-170	1.02 0.98 1.05	R13c R13a*R2*R21*R22 R13b*R11*R2*R21*R22	 2,6-8 4-8
VM/ESA 1.5 (370)	VM/ESA 2.1.0	9221-120 9221-170	0.96 1.03	(1/R13c)*R13a*R2*R21* R22 (1/R13c)*R13b*R11* R2*R21*R22	2,6-8 4-8
VM/SP HPO 5	VM/ESA 2.1.0	UP, ¬4381 MP, ¬4381	0.99 1.10	RHa*R2*R21*R22 RHb*R1E*R11* R2*R21*R22	4,5,7,8 3-5,7,8
VM/XA 2.0	VM/ESA 2.1.0		1.22	RX20*RX21*R1E* R11*R2*R21*R22	8
VM/XA 2.1	VM/ESA 2.1.0		1.19	RX21*R1E*R11* R2*R21*R22	8
VM/ESA 1.0 ESA	VM/ESA 2.1.0		1.15	R1E*R11*R2*R21*R22	8
VM/ESA 1.1	VM/ESA 2.1.0		1.10	R11*R2*R21*R22	8
VM/ESA 2	VM/ESA 2.1.0		1.09	R2*R21*R22	8
VM/ESA 2.1	VM/ESA 2.1.0		1.08	R21*R22	8
VM/ESA 2.2	VM/ESA 2.1.0		1.05	R22	8

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 79.

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 79 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 1.0.0 (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 2.1.0.

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 1.1.5 (370 Feature) ITRR is based (in part) on a comparison of VM/SP 5 to VM/ESA 1.0.0 (370 Feature), which showed an ITRR of 0.92. This comparison was done with 16MB of real memory. However, VM/ESA 1.0.0 (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 1.0.0 (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 2.1.0 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:
   - Use 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 1.0.0 (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.
8. This ITRR value depends strongly upon the fact that CMS is now shipped with most of its REXX execs and XEDIT macros compiled (see “Performance Improvements” on page 9). If these are already compiled on your system, divide the ITRR shown by 1.07.

Table 23 on page 76 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.

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/ESA 2.1 to VM/ESA 2.1.0. From Table 23 on page 76, the estimated ITRR for migrating from VM/ESA 2.1 to VM/ESA 2.1.0 is 1.08. Therefore, the estimated overall increase in system capacity is 1.30*1.08 = 1.40.

Table 23 on page 76 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 74.

Derivation and Supporting Data

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

The derivation column in Table 23 on page 76 shows how the stated ITR ratio was calculated. For example, the ITRR of 1.08 for migrating from VM/ESA 2.1 to VM/ESA 2.1.0 was calculated by multiplying the average ITRR for migrating from VM/ESA 2.1 to VM/ESA 2.2 (R21) by the average ITRR for migrating from VM/ESA 2.2 to VM/ESA 2.1.0 (R22): 1.03*1.05 = 1.08. R21 was calculated by averaging the ITRRs for VM measurement pairs 24 through 27 (see Table 24 on page 80). Likewise, R22 was calculated by averaging the ITRRs for VM measurement pairs 28 through 30.

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 2.1.0 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 23 on page 76 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 23 on page 76 are based on the following pairs of measurements:

Table 24 (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)
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	
avg							1.01	(R11)
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)

Table 24 (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 2.1 to VM/ESA 2.2: FS8F0R; Report 22								
24	S46E5505	S47E550A	9121-742	1GB/1GB	90	17	1.026	
25	S46E5202	S47E5201	9121-742	320MB/64MB[8]	90	20	1.037	
26	L26E186I	L27E186J	9121-480	224MB/32MB[8]	90	16	1.026	
27	H16E0302	H17E0303	9221-170	48MB/16MB[8]	90	15	1.026	
avg							1.03	(R21)
VM/ESA 2.2 to VM/ESA 2.1.0: FS8F0R; Report 210								
28	S47E550D	S48E5500	9121-742	1GB/1GB	90	18	1.042	
29	L27E1909	L28E190M	9121-480	256MB	90	16	1.070	
30	H17E0304	H18E0303	9221-170	64MB	90	15	1.038	
avg							1.05	(R22)
Note: The report numbers refer to the list of VM performance reports on page 74.								

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 reads and writes.

Symbol The symbol used to represent this release transition in Table 23 on page 76.

The FS7B0R, FS7F0R, or FS8F0R workloads (CMS-intensive, minidisks, remote users simulated by TPNS) were used for all comparisons except those involving VM/SP 6. For those comparisons, the FS7B0 workload was used (CMS-intensive, minidisks, local users simulated by the full screen internal driver (FSID) tool).

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

[8] These are the storage sizes used for the VM/ESA 1.2.1 measurements. For VM/ESA 1.2.2, 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.

New Functions

A number of the functional enhancements in VM/ESA 2.1.0 have performance implications. This section contains performance evaluation results for the following functions:

- POSIX
- DCE
- GCS TSLICE Option

POSIX

This section provides performance information and measurement results for the POSIX support provided by VM/ESA 2.1.0. The following topics are covered:

POSIX Initialization
POSIX Functions: CPU Usage
POSIX Functions: Real Storage Requirements
Shell Initialization and Termination
Shell Commands: CPU Usage
Shell Commands: Real Storage Requirements
Large File Performance
Byte File System Loading

POSIX Initialization

The POSIX environment is implicitly initialized in a virtual machine whenever the first POSIX-oriented request is issued. This request might be, for example, an OPENVM MOUNT request. Once the POSIX environment is initialized, it remains until the virtual machine is reset (IPL CMS).

The amount of time required to do POSIX initialization is relatively small. On an unconstrained 9021-900, the elapsed time was observed to be about 0.1 seconds and about 16 milliseconds of CPU time was required.

Perhaps the main performance implication of POSIX initialization is that about 640 additional non-shared pages are referenced. If there is no subsequent use of POSIX functions, these pages will typically be paged out to DASD. If there is subsequent use of POSIX functions, a subset of these pages will continue to be referenced and additional pages may be referenced as well (see "POSIX Functions: Real Storage Requirements" on page 86).

Once the POSIX environment has been initialized, there is a slight increase in resource requirements for subsequent execution of non-POSIX CMS work. For example, an instruction trace of an EXEC that copies a small file, XEDITs it, and then erases that file showed a 0.25% increase in virtual machine instructions executed and a 5 page increase in referenced non-shared pages when it was executed after POSIX initialization.

Because of these various additional resource requirements, we recommend that you not put POSIX-oriented commands such as OPENVM MOUNT in your PROFILE EXEC unless you will normally be using POSIX functions subsequent to starting CMS.

POSIX Functions: CPU Usage

CPU usage information was obtained for a selection of frequently used POSIX functions. Elapsed time, byte file system (BFS) server calls, and BFS DASD I/Os were also collected.

The data reported for each function represent an average of multiple (typically 50) loop iterations of that function. The functions were executed by a C program, while data collection was controlled by assembler routines called by that program. Elapsed times were obtained using the STCK instruction. User machine CPU times were obtained using diagnose code X′0C′. Byte file system server CPU times were obtained using the CP QUERY TIME command. BFS server statistics were collected using the QUERY FILEPOOL COUNTER command. All this information was collected immediately prior to and immediately following execution of each function loop.

The measurements were made on a non-dedicated 9021-900 during low usage conditions when contention from other system activity was minimal. The BFS server was dedicated during the measurement. Multiple measurement runs were obtained to verify repeatability. The results, shown in Table 25, are from a typical measurement run.

Table 25 (Page 1 of 2). Performance of individual POSIX calls on a 9021-900 processor

Function	Elapsed Time (μsec)	Total CPU Time (μsec)	User CPU TIme (μsec)	Server CPU Time (μsec)	BFS Calls	BFS I/Os
chmod()	7699	3476	1676	1800	2	2
close()	1506	1574	974	600	1	0
closedir()	1470	1600	1000	600	1	0
creat()	10777	5114	2114	3000	2	2
fcntl()	69	68	68	0	0	0
fstat()	116	112	112	0	0	0
fsync()	1720	1766	966	800	1	0
ftruncate()	8102	2416	1016	1400	1	2
getcwd()	1538	1790	990	800	1	0
getlogin()	12	12	12	0	0	0
getpid()	4	4	4	0	0	0
getppid()	4	4	4	0	0	0
link()	11222	6862	3062	3800	3	2
lseek()	66	64	64	0	0	0
mkdir()	11739	6474	2874	3600	3	2
mkfifo()	14124	6790	2990	3800	3	2
open()	10474	4858	2058	2800	2	2
opendir()	1785	1762	1162	600	1	0
pipe()	465	462	462	0	0	0
read()	361	152	152	0	0	0
readdir()	324	236	36	200	0.02	0
readlink()	233	232	232	0	0	0
rename()	12941	8086	4086	4000	4	2
rmdir()	9065	4404	1804	2600	2	2
sigprocmask()	7	7	7	0	0	0
stat()	251	248	248	0	0	0
symlink()	12677	6770	2970	3800	3	2
time()	2	2	2	0	0	0

Table 25 (Page 2 of 2). Performance of individual POSIX calls on a 9021-900 processor

Function	Elapsed Time (µsec)	Total CPU Time (µsec)	User CPU TIme (µsec)	Server CPU Time (µsec)	BFS Calls	BFS I/Os
times()	197	38	38	0	0	0
unlink()	10104	5624	2624	3000	3	2
utime()	7226	2712	1312	1400	1	2
write()	363	152	152	0	0	0

Notes:

1. All times are in microseconds.
2. The measurement accuracy of server CPU time is limited by the accuracy of the QUERY TIME command (hundredths of a second).
3. For getpid(), getppid(), getlogon(), time(), times(), and sigprocmask(), 1000 loop iterations were used. Fifty loop iterations were used for all other measured functions.
4. All files, directories, links, etc, were created in the root directory.
5. The performance of a given function can vary depending upon how it is invoked and other conditions. The following list provides qualifying information for the measured cases:

chmod()	grant all permissions to a file
close()	the file is empty
creat()	file does not exist
fcntl()	obtain file status and file access mode flags
fsync()	no modified data to be forced out
ftruncate()	zero-length file truncated to 0 bytes
getcwd()	root directory
lseek()	position to beginning of a zero-length file
open()	file does not exist
read()	100 bytes, successive reads are sequential
sigprocmask()	SIG_BLOCK
stat()	file exists and was recently referenced
write()	100 bytes, successive writes are sequential

The results indicate that POSIX functions that do not need to make any trips to the BFS server require much less CPU time than those that do. Also, CPU time is roughly proportional to the number of BFS server calls that are required.

For the cases where the measured function required one or more trips to the BFS server but there were no server I/Os, the elapsed time is often somewhat lower than total CPU time. This is because there is a slight amount of overlap between processor usage in the user machine and processor usage in the BFS server machine (the measurements were obtained on a 6-way processor).

The counts and timings provided by the QUERY FILEPOOL COUNTER command are helpful for understanding BFS server activity. Table 26 shows sample counter data for one of the measured POSIX functions.

Table 26. Sample Q FILEPOOL COUNTER data: open()	
Count/ Open	**Counter Description**
1.00	Byte File Lookup Requests
1.00	Byte File Open File New With Intent Write Requests
0.02	Query File Pool Requests
2.00	Total Byte File File Pool Requests
2.02	Total File Pool Requests
8.04	File Pool Request Service Time (msec)
2.02	Local File Pool Requests
19.14	SAC Calls
1.00	LUW Rollbacks
2.02	Begin LUWs
8.70	Agent Holding Time (msec)
2.40	Log Blocks Written
2.40	Total DASD Block Transfers
2.00	BIO Requests to Write Log Blocks
2.00	Total BIO Requests
6.08	Total BIO Request Time (msec)
2.00	I/O Requests to Write Log Blocks
2.00	Total I/O Requests

These data were obtained by issuing the QUERY FILEPOOL COUNTER command immediately before and after executing the open() function loop, subtracting the results, and dividing by 50 (the number of loop iterations).

The results show that two BFS requests are required to implement the open. (The purpose of the lookup request is to determine whether the file already exists.) You can also see that the open resulted in two I/O requests (a forced write to each of the two logs), those I/O requests took a total of 6.08 msec to complete, and that this accounts for most of the 8.04 msec it took for the two BFS requests to be handled by the BFS server (File Pool Request Service Time).

POSIX Functions: Real Storage Requirements

Instruction traces of the user virtual machine were collected for a subset of the POSIX functions listed in Table 25. This information was used to determine the number of unique pages that were referenced during the execution of each traced function and during the combined execution of all the traced functions. Each referenced page was classified as non-shared or shared based upon whether there is one such page per CMS user or whether there is just one instance of that page that is shared among all CMS users. The results are summarized in Table 27.

Table 27. Real storage requirements of selected POSIX calls

Function	Unique Non-Shared Pages	Unique Shared Pages
close()	77	81
fsync()	34	23
getpid()	18	3
open()	88	92
read()	40	30
sigprocmask()	20	4
write()	30	30
all	130	106

The count of non-shared pages is more important than the count of shared pages because each user doing that function requires a separate set of these pages.

All of the shared pages are from the CMS saved system.

The unique page reference counts listed in the table are for the user virtual machine. Those POSIX functions that require calls to the byte file system server will also cause pages in the BFS server to be referenced. Those page references, however, are less important because there is just one set of those pages in a server virtual machine that is (potentially) servicing many end users.

POSIX functions that use the byte file system tend to touch more non-shared pages than do analogous shared file system calls. An example CMS command that uses open, close, read, and write SFS calls was found to reference 40 non-shared pages.

Shell Initialization and Termination

Invocation of the OPENVM SHELL command causes the POSIX shell to be initialized. You remain in the shell environment until you specify ″exit″.

Shell initialization references about 1030 pages in the user virtual machine (in addition to the 640 pages that are referenced during POSIX initialization). A subset of these pages continue to be referenced when shell commands are being executed (see “Shell Commands: Real Storage Requirements” on page 89). About 980 of these pages are released on exit from the shell environment. These figures were determined through use of the CP INDICATE USER * EXPANDED command. The sum of the user′s resident, expanded storage, and DASD pages was used as a measure of how many unique pages were referenced.

Total CPU time required to initialize and exit the shell was measured using the CP QUERY TIME command. On a 9021-900, shell initialization CPU time was about 0.51 seconds, while shell termination CPU time was about 0.03 seconds.

Shell Commands: CPU Usage

CPU usage information was obtained for a selection of frequently used POSIX shell commands. Elapsed time, BFS server calls, and BFS DASD I/Os were also collected.

The measured commands were collected into a shell script. Each measured command was immediately preceded and immediately followed by a data collection program, invoked by the ″cms″ shell command. Elapsed times were obtained using the STCK instruction. User machine CPU times were obtained using diagnose code X′0C′. Byte file system server CPU times were obtained using the CP QUERY TIME command. BFS server statistics were collected using the QUERY FILEPOOL COUNTER command.

The data reported for each command are based upon a single execution of that command. The result shown have been adjusted to subtract out the elapsed time and CPU time required to collect the data.

The measurements were made on a non-dedicated 9021-900 during low usage conditions when contention from other system activity was minimal. The BFS server was dedicated during the measurement. Multiple measurement runs were obtained to verify repeatability. The results, shown in Table 28, are from a typical measurement run.

Table 28. Performance of POSIX shell commands on a 9021-900 processor

Command	Elapsed Time (sec)	Total CPU Time (sec)	User CPU Time (sec)	Server CPU Time (sec)	BFS Calls	BFS I/Os
cat_20	0.21	0.12	0.11	0.01	12	19
cat_100	0.44	0.30	0.26	0.04	16	37
cd	0.03	0.01	0.01	0.00	1	9
chmod	0.08	0.06	0.05	0.01	8	10
cp_1	0.19	0.09	0.07	0.02	16	26
cp_20	0.30	0.09	0.07	0.02	14	23
cp_100	0.42	0.11	0.08	0.03	18	44
c89_null	4.05	0.57	0.53	0.04	37	41
c89_72k	5.43	2.13	2.05	0.08	78	55
diff	0.45	0.36	0.34	0.02	25	17
echo	0.01	0.02	0.01	0.01	1	0
grep	0.14	0.07	0.06	0.01	11	15
ls	0.18	0.13	0.10	0.03	25	34
mkdir	0.16	0.07	0.05	0.02	11	29
mv	0.15	0.06	0.05	0.01	11	26
ps	0.09	0.06	0.05	0.01	12	13
pwd	0.01	0.01	0.00	0.01	1	3
rm	0.13	0.07	0.05	0.02	9	26
rmdir	0.14	0.07	0.05	0.02	8	37
sort_1	0.36	0.12	0.08	0.04	27	40
sort_20	0.54	0.43	0.40	0.03	26	30
sort_100	6.14	5.69	5.63	0.06	31	77

Notes:

1. A command name that is suffixed with ″_nn″ means that the file involved is nn Kbytes in size. For example, ″cp_20″ means that the cp command copies a 20KB file.

2. ″c89_null″ means that the c89 command is used to compile a null (4 lines) C program. ″c89_72″ means that the c89 command is used to compile a 72KB (2050 lines) C program.

For all but the c89 cases, BFS I/Os represent all the DASD I/Os that occurred during that command′s execution. The c89 command caused additional DASD I/Os to occur because it used a number of CMS files during the compilation. For example, the included header files were in an SFS directory that resided in a separate SFS file pool.

The total CPU time required to run a shell command (for example, “rm”) is typically much higher than the CPU time required to run an analogous CMS command (for example, “erase”). A major contributing factor is that shell commands must be run in an asynchronous, multitasking environment whereas CMS commands run serially on the command thread. In spite of the increased CPU usage, most shell commands are quite responsive. Most of the commands in Table 28 show subsecond response time.

Shell Commands: Real Storage Requirements

Instruction traces of the user virtual machine were collected for two example shell commands — ″ls″ and ″rm″. The ″ls″ command listed the (16) files in the current directory. The ″rm″ command removed a small file from the current directory. These traces were used to determine the number of unique pages that were referenced during the execution of each traced command and during the combined execution of both traced commands. Each referenced page was classified as non-shared or shared based upon whether there is one such page per CMS user or whether there is just one instance of that page that is shared among all CMS users (by being in a shared segment or saved system). The results are summarized in Table 29.

Table 29. Real storage requirements of two example shell commands

Command	Unique Non-Shared Pages	Unique Shared Pages
rm	572	312
ls	571	311
both	709	319

The count of non-shared pages is more important than the count of shared pages because each user doing that function requires a separate set of non-shared pages. All of the shared pages are from the CMS saved system.

Such a small sample of commands can only give a very general idea of the real storage requirements associated with running shell commands. The results suggest that most shell commands tend to reference the same shared pages but they reference somewhat different subsets of the non-shared pages.

Large File Performance

A series of single-user measurements was obtained to explore the performance of sequentially writing and reading large BFS files using the write() and read() functions. File size and the number of bytes transferred per request were varied.

Elapsed times were obtained by an assembler routine that used the STCK instruction. User machine CPU times were obtained using the clock() function. Byte file system server CPU times were obtained using the CP QUERY TIME command. BFS server statistics were collected using the QUERY FILEPOOL COUNTER command.

The measurements were taken on a non-dedicated 9021-900 during low usage conditions when contention from other system activity was minimal. The BFS server was dedicated during the measurements. Multiple measurement runs were obtained to verify repeatability. The results, shown in Table 30, are from a typical measurement run.

Table 30. Large file performance on a 9021-900 processor

Case	Elapsed Time (sec)	MB/ Second	Total CPU Time (sec)	User CPU Time (sec)	Server CPU Time (sec)	BFS Calls	BFS I/Os
write, 0.5MB, 128 bytes/request	0.96	0.52	0.63	0.55	0.08	33	88
write, 0.5MB, 4KB/request	0.46	1.08	0.13	0.05	0.08	33	88
write, 0.5MB, 64KB/request	0.44	1.13	0.10	0.03	0.07	33	87
write, 2MB, 128 bytes/request	3.88	0.52	2.51	2.23	0.28	110	329
write, 2MB, 4KB/request	1.82	1.10	0.46	0.18	0.28	110	331
write, 2MB, 64KB/request	1.73	1.16	0.41	0.13	0.28	110	330
read, 0.5MB, 128 bytes/request	0.76	0.66	0.58	0.55	0.03	29	32
read, 0.5MB, 4KB/request	0.25	2.00	0.08	0.05	0.03	29	32
read, 0.5MB, 64KB/request	0.25	2.02	0.06	0.03	0.03	29	32
read, 2MB, 128 bytes/request	3.08	0.65	2.37	2.26	0.11	158	120
read, 2MB, 4KB/request	1.06	1.89	0.31	0.21	0.10	158	120
read, 2MB, 64KB/request	0.97	2.06	0.26	0.16	0.10	158	120

Most of the BFS calls are write or read requests. For example, in the first case listed in the table, 26 of the 33 BFS calls are write requests. The byte file system reads ahead and writes behind up to 5 4KB blocks at a time.[9] This is reflected in these results. For example, the first case involves writing a 0.5MB file, which is 128 4KB blocks. 128 blocks / 26 BFS writes = 4.92 blocks per BFS write.

In the read cases, the number of BFS I/Os is approximately equal to the number of BFS calls. This reflects the fact that most of these calls are read requests and most of those requests are satisfied by the use of one multi-block I/O request. In the write cases, the number of BFS I/Os is approximately 3 times higher than the number of write requests. This reflects the fact that each write request typically results in one multi-block I/O request to write the file data and one forced write to each of the two BFS file pool logs.

[9] This is analogous to, but different from, the read-ahead, write-behind mechanism used by the CMS file system (minidisks and SFS).

User machine CPU time decreases as the number of bytes per write or read request increases, reflecting the increased efficiency achieved by having fewer function calls to process. Server machine CPU time is essentially independent of the number of bytes per write or read request because this has little effect on the number of BFS calls that are made or the number of bytes requested per BFS call.

Byte File System Loading

A series of measurements was collected to assess the performance characteristics of a BFS server as a function of increased loading.

The load was applied by 1 to 8 disconnected user virtual machines running concurrently. Each user machine executed the same workload. There was no delay (think time) between requests so the user machines represent batch activity. The workload consisted of writing, reading, and erasing large files (13 0.5MB files, 14 1MB files, and 15 2MB files). Each user machine had a separate working directory and worked with its own set of files. However, all of these files resided in one byte file system in the same BFS server.

The measurement interval was from first user start until last user end. All users started and ended their work within a few seconds of each other. Monitor data were collected at 10-second intervals during this period and later reduced by VMPRF. QUERY FILEPOOL COUNTER report output was collected before and after the measurement interval.

The BFS server data (storage group 1 and storage group 2) were spread across 10 3390-2 volumes. Two of these volumes also contained a log minidisk. Five volumes (with one log minidisk) were behind one 3990-3 control unit, while the other five volumes (with the other log minidisk) were behind a separate 3990-3 control unit.

Workload: Reading and Writing Large BFS Files

Hardware Configuration

Processor model:	9121-480
Processors used:	2
Storage:	
Real:	256MB
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				
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2			10 R[10]	

[10] DASD fast write was enabled for one of the measurements (see discussion).

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 Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
BFSERV1	1	BFS	64MB/XC	1500	1300	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1-8	Users	32MB/XC	100		

Measurement Discussion: Megabytes per second (read or written) was selected as the measure of overall throughput. Figure 9 shows a plot of throughput as a function of the number of concurrent user virtual machines. Table 31 summarizes the performance information that was collected at each measured number of users.

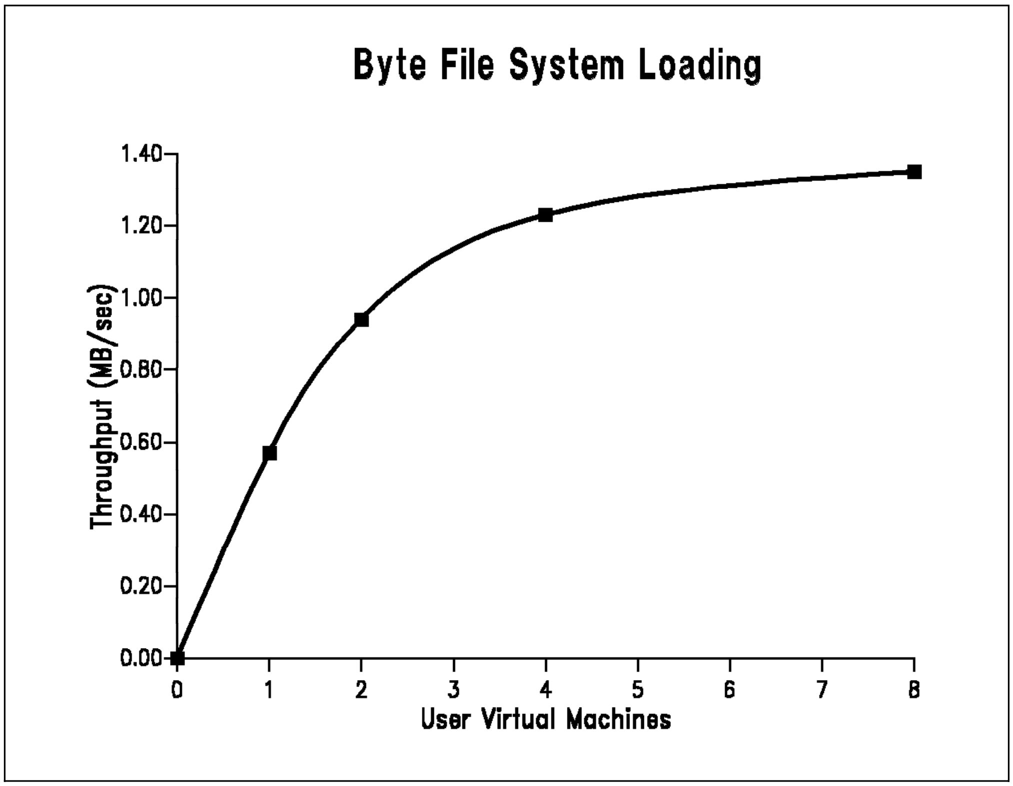

Figure 9. BFS throughput vs number of users on a 9121-480

The results show that the BFS server was able to support a high level of concurrency. The measured configuration achieved a maximum throughput exceeding 1.4 MB/sec. As discussed below, the I/O subsystem was the limiting factor in this case.

Table 31. BFS performance vs number of batch users on a 9121-480

Batch User Machines	1	2	4	8
Elapsed Time (sec)	177	214	328	597
Megabytes Transferred	101	202	404	808
MB/sec	0.57	0.94	1.23	1.35
Files Processed	84	168	336	672
Files/sec/User	0.475	0.393	0.256	0.141
Seconds/File	2.11	2.54	3.91	7.11
CPU-Seconds/File	1.03	1.06	1.08	1.20
CPU Utilization	24.5	41.7	55.1	60.2
BFS Server Pct Busy	13.9	24.3	31.4	35.6
Highest Channel Busy	8.9	9.5	9.4	9.0
Log DASD (MDSK09):				
IOs/sec	17.3	29.4	38.3	40.5
Percent Busy	23.1	38.4	45.4	51.2
Service Time (msec)	13.3	13.0	11.9	12.7
Response Time (msec)	13.3	15.3	14.5	15.2
Avg Non-log DASD:				
IOs/sec	2.0	3.1	4.1	4.5
Percent Busy	2.8	5.8	8.5	9.2
Service Time (msec)	14.4	18.9	20.9	23.4
Response Time (msec)	14.4	20.4	33.4	51.5

Note: All CPU, channel, and DASD performance data are from VMPRF.

Notes:

1. Seconds/File is 1/(Files/sec/User). It is the average elapsed time required to read or write a file. The average file is 1.2MB in size.
2. CPU-Seconds/File is (CPU Time)/(Files Processed), where CPU Time is (Elapsed Time)*2*(CPU Utilization)/100. CPU Utilization is multiplied by 2 because there are 2 processors on the 9121-480.
3. BFS Server Pct Busy is (total CPU-seconds used by the BFS server)/(Elapsed Time). It represents the percentage of the time that the BFS server is running on a processor. Since the BFS server can only run on one processor at a time, this cannot exceed 100% and thus represents a potential limiting factor.
4. The two BFS logs are on MDSK01 and MDSK09. Because the BFS file pool was implemented symmetrically across two control units and because log writes are done to both logs, the performance results for each log DASD are nearly identical. Results for MDSK09 are shown here.
5. The performance characteristics of the remaining 8 non-log DASD in the BFS file pool are averaged and shown under Avg Non-log DASD.
6. Service Time is the average amount of time it takes to complete the DASD I/O request once it has been initiated.
7. Response Time is Service Time plus the amount of time spent waiting before the I/O can be initiated because the path to the required DASD volume is busy.

For the measured configuration, the I/O subsystem was the limiting factor that determined maximum throughput. More specifically, contention at the two control units appears to have been the limiting factor. This is indicated by the Avg Non-log DASD results. Service Time remains at reasonable levels while progressing from 1 to 8 users, while Response Time increases greatly. The moderate DASD service times and low DASD utilizations show that there is no problem doing the I/Os once the path to the DASD is clear. The very low channel utilizations show that channel contention is not a problem.

The primary source of control unit contention was from the very high I/O rates to the BFS logs. This meant that whenever an I/O was to be started to a non-log DASD, it was likely that the control unit was busy handling a log I/O request. The DASD response time of the log DASD was less affected at high loadings because the other 4 DASD behind the same control unit had relatively low I/O rates.

There are, of course, many other potential limiting factors. Which factor serves to limit throughput depends upon the relative capacities of the various hardware components in the configuration. For these measurements, the overall capacity of the 480 processor was not limiting, as evidenced by the fact that CPU Utilization only reached 60%. The CPU utilization in the BFS server is another potential limiting factor, but BFS Server Pct Busy only went up to 36%.

The fact that the limiting factor can change depending upon the configuration is illustrated by an additional 8-user measurement that was done with DASD fast write (DFW) enabled for the 10 BFS file pool DASD (see Table 32). The presence of write caching greatly improved the I/O subsystem, allowing the throughput rate to rise from 1.35 to 1.95 MB/sec. This new configuration is limited by the 9121-480′s processing capacity, as evidenced by the 86% CPU utilization.

At high levels of contention, deadlock conditions can occasionally occur in the BFS server. When the BFS server identifies a deadlock condition, it breaks the deadlock by failing one of the participating requests. This occurred during an attempted measurement at 12 users and the message ″EDC5116I Resource deadlock avoided″ was displayed on the affected user virtual machine′s console.

Table 32. The effect of DASD fast write on BFS throughput

DASD Fast Write **Batch User Machines**	**NO** **8**	**YES** **8**
Elapsed Time (sec)	597	414
Megabytes Transferred	808	808
MB/sec	1.35	1.95
Files Processed	672	672
Files/sec/User	0.141	0.203
Seconds/File	7.11	4.93
CPU-Seconds/File	1.20	1.06
CPU Utilization	60.2	86.3
BFS Server Pct Busy	35.6	49.8
Highest Channel Busy	9.0	13.2
Log DASD (MDSK09):		
IOs/sec	40.5	58.2
Percent Busy	51.2	19.9
Service Time (msec)	12.7	3.4
Response Time (msec)	15.2	3.8
Avg Non-log DASD:		
IOs/sec	4.5	6.2
Percent Busy	9.2	5.6
Service Time (msec)	23.4	8.9
Response Time (msec)	51.5	14.6

Note: All CPU, channel, and DASD performance data are from VMPRF.

DCE

This section provides performance measurement results for IBM OpenEdition Distributed Computing Environment for VM/ESA (VM DCE). The following topics are covered:

- Idling Overhead
- Single Thread RPC Measurements
- RPC Throughput Measurements

All measurement results provided in this section were made with TCP/IP for VM Version 2 Release 3 and a pre-ship version of VM DCE running on VM/ESA 2.1.0.

Idling Overhead

The DCECORE virtual machine and VM DCE application servers set up various timer events and take timer interrupts as part of carrying out their functions. As a result, there is a small amount of overhead associated with these server virtual machines once they are started, even when they are not in use.

Three measurements were taken on a dedicated 9121-320 system to quantify this idling overhead. Monitor records were collected for each measurement and later reduced by VMPRF. The first measurement was a base measurement of an idle VM/ESA 2.1.0 system taken before any DCE-related servers were started. The SYSTEM_SUMMARY_BY_TIME report showed that the system was 1.8% busy. This was due to CP functions (such as scheduling and monitor), MONWRITE, and TCP/IP.

The second measurement was taken about a minute after starting the DCECORE virtual machine. The system utilization rose to 2.1%, indicating that the DCECORE virtual machine was using 0.3% of the processor.

A VM DCE server application was then started and the third measurement was taken. The system utilization remained at 2.1%. The USER_RESOURCE_UTIL report showed a slight amount of processor usage in the server virtual machine, but it was apparently not enough to affect the reported total system processor usage, which is shown with a precision of 0.1%. Each application server virtual machine has its own timer activity, so several such servers could have a measurable effect on total system idling overhead.

Since the frequency of timer interrupts handled by the DCECORE virtual machine and DCE application servers is independent of processor speed, you can expect the DCE-related idling overhead, as a percentage of total processor capacity, to be higher than 0.3% on slower processors and lower than 0.3% on faster processors.

Single Thread RPC Measurements

Single thread remote procedure call measurements were obtained for 24 different RPC types on 3 different hardware configurations.

The RPCs were executed by an internal RPC performance driver application. The client side of the application was executed on an AIX system running on an RS/6000. The server side was executed in a VM/ESA virtual machine running on an ES/9000 system. The client and server machines were in the same DCE cell.

The AIX and VM/ESA systems were connected through a 16 megabit IBM Token Ring. During the measurements, the AIX and VM/ESA systems were dedicated. The token ring was not dedicated so the reported response times are somewhat influenced by the presence of other activity on the LAN. However, repeat measurements and measurements obtained during low usage hours indicate that extraneous activity on the LAN did not have an appreciable effect on the results.

The 24 measured RPC cases consisted of 4 different sizes, using each of the 6 available RPC protection levels. For each such case, the measurement consisted of one AIX client thread sending consecutive RPCs (no think time) to the application server on VM/ESA for 2 minutes. For each RPC, the client sent the requested number of bytes to the server and waited for the server′s response. The server echoed the same amount of data back to the client. For both the send and the reply, the data were in the form of a single binary argument value.

RPC counts and RPC response time data were obtained from the performance driver. Monitor records were collected on the VM/ESA system at 6-second intervals and reduced by VMPRF. The CPU usage data in the SYSTEM_SUMMARY_BY_TIME and USER_RESOURCE_UTIL reports were used, in conjunction with the RPC counts, to calculate ES/9000 CPU usage per RPC.

A set of RPC measurements was collected on three different configurations.

Table 33. Measured configurations

	Config 1	**Config 2**	**Config 3**
AIX system (RS/6000 model)	220	520	520
AIX release level	3.2.4	3.2.5	3.2.5
VM/ESA system	9121-320	9121-621[11]	9121-621[11]
VM/ESA system processors	1	2	2
host/LAN connection	3172-1	3172-3	OSA-1[12]
For results, see ..	Table 34	Table 35	Table 36

[11] For config 2, this system was configured by physically partitioning a 9121-742 (4-way) or by varying two of the processors offline from the hardware configuration screen. For config 3, this system was configured by bringing up a single image 9121-742 and then using the CP VARY PROCESSOR command to vary two of the alternate processors offline. This was done to expedite switching between 3172-3 and OSA-1 LAN connectivity. The performance differences between these methods of configuring a “9121-621” are small and can be ignored when making comparisons.

[12] IBM S/390 Open Systems Adapter

Table 34. Single thread RPC performance: Config 1 (9121-320, 3172-1, RS/6000 220)

KB Sent/ Returned	Protection Level	Response time		9121-320 CPU time per RPC			
		Average	Low	Total	Server	TCP/IP	Other
0	none	20.5	19.2	13.5	9.1	3.6	0.8
1	none	39.9	24.7	14.2	9.3	4.0	0.9
4	none	40.5	29.7	14.7	9.1	4.4	1.2
32	none	171.0	146.2	90.3	55.5	28.4	6.4
0	connect	21.0	19.8	13.8	9.3	3.7	0.8
1	connect	40.2	24.6	14.5	9.4	4.0	1.1
4	connect	40.7	29.4	15.8	9.5	4.4	1.9
32	connect	171.1	148.0	91.5	56.8	28.4	6.3
0	call	21.0	19.7	13.8	9.3	3.7	0.8
1	call	40.2	24.1	14.5	9.4	3.7	1.4
4	call	40.6	30.0	15.0	9.5	4.4	1.1
32	call	171.5	147.6	91.0	57.0	28.5	5.5
0	pkt	21.0	19.7	13.8	9.2	3.7	0.9
1	pkt	40.2	24.9	14.5	9.4	4.0	1.1
4	pkt	40.6	29.8	15.0	9.5	4.4	1.1
32	pkt	170.5	148.5	91.1	56.7	28.3	6.1
0	pkt_integ	22.0	20.6	14.4	9.7	3.7	1.0
1	pkt_integ	40.3	27.6	16.6	11.4	3.7	1.5
4	pkt_integ	51.2	39.4	22.2	16.2	4.3	1.7
32	pkt_integ	218.5	198.5	146.8	110.5	29.0	7.3
0	pkt_privacy	23.2	21.8	15.1	10.5	3.7	0.9
1	pkt_privacy	50.5	36.2	22.4	16.8	3.8	1.8
4	pkt_privacy	83.3	68.6	43.2	36.7	4.8	1.7
32	pkt_privacy	400.5	383.8	307.9	265.8	29.9	12.2

Note: All times are in milliseconds. All CPU results are from VMPRF.

The protection levels have the following meanings:

none No protection.

connect Provide protection when the client connects with the server.

call Provide protection when the server receives the request.

pkt Ensure that all data received is from the expected client.

pkt_integ Ensure that none of the data transferred between client and server has been modified.

pkt_privacy All of the above protection plus encryption.

Like all response times, RPC response times are determined by the combination of many variables such as processor speeds, network capacity, and resource contention. Because of this, the RPC response time results shown in this and the following tables should be viewed as illustrative examples.

Because the token ring was not dedicated, the average response times are increased to some extent by the presence of extraneous activity on the LAN. This extraneous activity has little effect on low response time, which typically occurs when token ring contention is not present at that point in time. If the token ring had been dedicated, it is expected that the average response times would fall between the low and average response time values shown in the results tables.

Total CPU time per RPC is based on the Pct Busy column in the SYSTEM_SUMMARY_BY_TIME VMPRF report. This time includes all CPU usage in the system, including CP system time that cannot logically be attributed to any given user. Server and TCP/IP time per RPC are based on the Total CPU Seconds column in the USER_RESOURCE_UTIL report. This includes all virtual CPU time consumed by that virtual machine plus all CP CPU time that is used to satisfy its requests for CP services. Other CPU time per RPC is calculated as Total - (Server + TCP/IP). This component basically corresponds to the non-DCE and DCE idling overhead discussed in "Idling Overhead" on page 96. Most of it is CP system time and is not related to the RPC activity being measured.

The results show that the performance of the first four protection levels is about the same and is essentially independent of RPC size. In contrast to this, the performance of pkt_integ and pkt_privacy is very size-dependent. For the null RPC case (0KB), response times and CPU usage are not much higher for pkt_integ and pkt_privacy as compared to the first four protection levels. The difference grows rapidly, however, as RPC size increases.

As might be expected, response time and ES/9000 CPU usage increase with increasing RPC size.

VM DCE uses a maximum packet size of 4274 bytes. This is hardcoded and is not a tuning variable. For the first four DCE protection levels (those whose cost is independent of RPC request size), CPU time per RPC is roughly proportional to the number of RPC packets required to contain the request. This is not shown very well in the results table because it lacks intermediate sizes between 4KB and 32KB, but this has been verified by additional measurements at 8KB and 16KB.

The packet size used by the lower level transport layer is typically smaller. When that is the case, these packets are combined into RPC packets and fragmented from RPC packets by TCP/IP. This is the main reason why CPU time per RPC in TCP/IP increases when going from 0KB to 1KB to 4KB RPCs. For the measurements in this report, the transport layer packet size was 1500 bytes (plus header length).[13]

All of the RPC results in this report are for non-idempotent[14] requests. Additional measurements (not shown) yielded equivalent results for a corresponding set of RPCs that were declared as being idempotent. The results were equivalent because the think time used (zero) was short enough that separate acknowledgements to the server were not required so the idempotent optimization did not apply.[15]

---

[13] This is specified in the TCP/IP configuration file (the default name is PROFILE TCPIP). See "Tuning Performance" in *TCP/IP Version 2 Release 2 for VM: Planning and Customization* for performance considerations.

[14] A non-idempotent request must execute either once, partially, or not at all. An idempotent request can safely be done more than once. That is, even if it is done more than once, it yields the same results and produces no undesirable side effects. Each different RPC call type is declared as being idempotent or not in the interface definition language (IDL) file.

[15] If the RPC is non-idempotent, the RPC runtime in the client must send an acknowledgement to the server that it has successfully received the RPC response. If the client thread sends the server another RPC within 3 seconds, this subsequent RPC serves as the acknowledgement and there is no overhead to handle a separate acknowledgement. Consequently, an

Because these are single thread measurements and because the measured RPCs do not incur any I/O or other delays while they are being handled in the VM/ESA system, the total ES/9000 CPU time per RPC shows approximately how much of the observed average response time delay is in the server system. Consider the first RPC case (0 KB, protection level of none) as an example. About 13.5 msec of the 20.5 msec response time is spent in the 9121-320 server system. The remaining 7.0 msec represents time spent in the RS/6000 model 220 client system handling its half of the RPC processing, along with transmission latency in the token ring and in the 3172-1 control unit.

The measurements in this report are for the case where the bytes are transmitted as a single binary argument value. CPU usage per RPC would have only been slightly higher if the bytes had been sent as multiple binary argument values. The presence of arguments of a data type that requires conversion (such as character or floating point) will increase RPC CPU usage more significantly.

The measurements in this report are for the case where client and server both reside in the same DCE cell. Similar results can be expected for the case where the hardware configuration remains the same but the two nodes are configured in two separate DCE cells. The reason for this is that once the binding between client and server is complete, subsequent RPCs flow directly between client and server in the same manner regardless of what cells they reside in.

The measurements in this report are for the case where the server application resides in a VM/ESA system, while the client is on another node elsewhere in the network. The processing required to handle an RPC request on the client side is quite similar to the processing that is required on the server side. This has been confirmed by additional measurements (not shown) where the client side was run on the measured VM/ESA system and the server was on the AIX system.

idempotent RPC performs no better than an equivalent non-idempotent one if the client thread's RPCs are separated by less than 3 seconds. However, when the spacing is more than 3 seconds, an idempotent RPC performs better because the separate acknowledgement is avoided.

Table 35. Single thread RPC performance: Config 2 (9121-621, 3172-3, RS/6000 520)

KB Sent/ Returned	Protection Level	Response time		9121-621 CPU time per RPC			
		Average	Low	Total	Server	TCP/IP	Other
0	none	19.9	14.4	9.0	6.2	2.2	0.6
1	none	20.0	15.2	9.0	6.2	2.2	0.6
4	none	20.6	17.2	9.4	6.2	2.4	0.8
32	none	91.3	84.3	60.3	39.6	17.5	3.2
0	connect	20.0	13.8	9.0	6.4	2.2	0.4
1	connect	20.1	15.5	9.1	6.2	2.2	0.7
4	connect	21.1	19.0	9.6	6.4	2.5	0.7
32	connect	91.7	84.8	61.1	40.5	17.6	3.0
0	call	20.0	13.7	9.0	6.4	2.0	0.6
1	call	20.2	15.7	9.1	6.3	2.2	0.6
4	call	21.1	19.0	9.5	6.4	2.5	0.6
32	call	91.3	85.1	61.0	40.4	16.8	3.8
0	pkt	19.8	13.5	9.0	6.3	2.2	0.5
1	pkt	20.0	15.2	9.2	6.4	2.2	0.6
4	pkt	21.2	18.9	9.6	6.4	2.5	0.7
32	pkt	92.6	85.5	61.7	40.9	17.7	3.1
0	pkt_integ	20.0	14.1	9.5	6.7	2.2	0.6
1	pkt_integ	20.1	17.5	10.6	7.8	2.2	0.6
4	pkt_integ	30.3	27.3	14.7	11.1	2.5	1.1
32	pkt_integ	135.2	118.8	98.6	77.4	17.9	3.3
0	pkt_privacy	20.0	15.4	9.9	7.0	2.2	0.7
1	pkt_privacy	30.2	26.4	14.9	11.8	2.3	0.8
4	pkt_privacy	62.7	58.6	29.9	25.6	2.6	1.7
32	pkt_privacy	343.0	332.4	212.9	188.0	17.1	7.8

Note: All times are in milliseconds. All CPU results are from VMPRF.

Compared to the config 1 results shown in Table 34, these config 2 results show lower response times and ES/9000 CPU usage. The CPU times per RPC are 30% to 37% lower. This is in proportion to the speed difference between the 9121-320 processor and one of the 9121-621 processors.

For most of the RPC cases, the response time decreases exceed the CPU time decreases. This is due to the other configuration differences (3172 model and RS/6000 model). The fact that the 9121-621 is a 2-way has little significance for these single thread measurements because, for the most part, only one processor at a time is being used.

Table 36. Single thread RPC performance: Config 3 (9121-621, OSA-1, RS/6000 520)

KB Sent/ Returned	Protect Level	Response time		9121-621 CPU time per RPC			
		Average	Low	Total	Server	TCP/IP	Other
0	none	12.2	11.5	8.9	6.1	2.2	0.6
1	none	13.3	12.7	9.0	6.3	2.1	0.6
4	none	16.9	15.6	9.7	6.2	2.7	0.8
32	none	68.3	64.7	63.8	40.5	19.9	3.4
0	connect	12.5	11.8	9.0	6.3	2.1	0.6
1	connect	13.8	12.9	9.1	6.3	2.2	0.6
4	connect	17.0	15.9	9.8	6.4	2.7	0.7
32	connect	68.9	65.4	64.3	40.8	20.1	3.4
0	call	12.5	11.8	9.0	6.3	2.1	0.6
1	call	13.7	11.4	9.1	6.3	2.2	0.6
4	call	17.0	15.9	9.8	6.4	2.7	0.7
32	call	68.9	65.2	64.3	40.7	20.1	3.5
0	pkt	12.5	11.3	9.0	6.3	2.1	0.6
1	pkt	13.6	12.9	9.1	6.3	2.2	0.6
4	pkt	17.1	15.9	9.8	6.3	2.7	0.8
32	pkt	69.7	63.0	64.4	40.6	19.7	4.1
0	pkt_integ	13.3	12.6	9.3	6.7	2.1	0.5
1	pkt_integ	16.3	15.5	10.7	7.8	2.2	0.7
4	pkt_integ	25.2	24.0	15.0	11.2	2.7	1.1
32	pkt_integ	113.4	97.7	103.4	77.4	20.8	5.2
0	pkt_privacy	14.4	13.6	9.8	7.1	2.2	0.5
1	pkt_privacy	25.3	24.3	15.2	11.8	2.3	1.1
4	pkt_privacy	58.0	55.6	30.6	25.6	2.9	2.1
32	pkt_privacy	331.2	324.0	217.6	186.8	19.2	11.6

Note: All times are in milliseconds. All CPU results are from VMPRF.

The only difference between config 2 and config 3 is that config 2 uses a 3172-3 for LAN connectivity, while config 3 uses OSA-1 (in TCP/IP passthrough mode). As a result, all differences beyond normal run variability should be attributable to the differences between these two host/LAN connectivity methods.

The OSA-1 results show CPU usages per RPC that are similar to the 3172-3 results. The results match within 2% for the 0KB and 1KB cases and are 2% to 6% higher for the 4KB and 32KB results. As might be expected, the CPU usage increases observed for the 4KB and 32KB cases are mostly in the TCP/IP virtual machine and presumably reflect differences in how the 3172-3 and OSA-1 interact with TCP/IP.

The OSA-1 results show significantly lower response times relative to the corresponding 3172-3 results. Most of the RPC cases have average response time decreases in the 20% to 40% range. The larger RPCs using pkt_integ and pkt_privacy are exceptions. They show smaller response time decreases because the response times for these cases are more dominated by CPU usage in the client and server systems.

RPC Throughput Measurements

A set of RPC throughput measurements was collected for each of the three configurations described in Table 33. A non-idempotent 1KB RPC with protection level of none was used for all of these measurements. The performance driver application that was used for the single thread measurements was also used for these measurements. For each configuration, the degree of loading on the ES/9000 server system was progressively increased by increasing the number of concurrent client threads until maximum capacity was achieved. As with the single thread measurements, there was no think time between RPC requests.

For each of the three configurations, a set of measurements was obtained for which all RPCs were directed to a single application server virtual machine. That server was started with 10 threads for processing incoming RPCs.

For the 9121-621 3172-3 measurements (config 2), an additional set of measurements was obtained using two application servers. For those measurements, half of the total load originated from one AIX RS/6000 client system and was directed to one server virtual machine, while an equivalent load originated from a second AIX RS/6000 client machine and was directed to a second 10-thread server virtual machine.[16]

Each measurement was 5 minutes long. The RPC throughput rate, RPC count, and RPC response time data were obtained from the performance driver. Monitor records were collected on the VM/ESA system at 6-second intervals and reduced by VMPRF. The CPU usage data in the SYSTEM_SUMMARY_BY_TIME and USER_RESOURCE_UTIL reports were used, in conjunction with the RPC counts, to calculate ES/9000 CPU usage per RPC. Data from these same reports were also used to determine average processor utilization and average server utilization. Average server utilization was calculated as:

100*(ServerTCPU/Servers)/ElapsedTime

ServerTCPU is total CPU time used by the application server virtual machine(s), in seconds.

Servers is the number of server virtual machines (1 or 2).

ElapsedTime is the measurement duration, in seconds.

Figure 10 shows a plot of RPC throughput as a function of the total number of concurrent client threads for each of the 3172-based measurement sequences. Table 37, Table 38, and Table 39 provide additional data for these measurements. Figure 12 compares the 3172-3 and the OSA-1 results for the 9121-621 single server case. The OSA-1 results are provided in Table 40.

[16] Two client machines were used for these 2 server measurements because RPC throughput would otherwise have become limited by the CPU capacity of the (RS/6000 model 520) client machine. One client machine could have been used if an RS/6000 model with sufficiently higher capacity had been available. The second client machine was an RS/6000 model 250 running AIX 3.2.5.

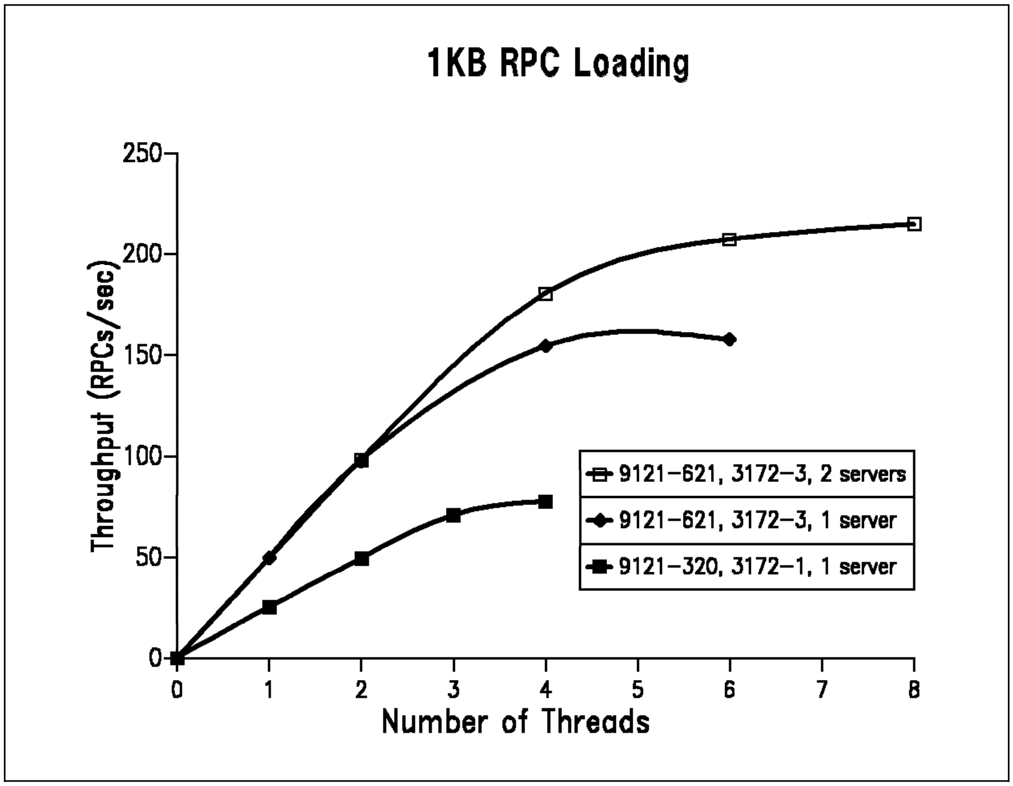

Figure 10. VM DCE throughput capacity on various ES/9000 processors

9121-320 Results

Table 37. 1KB RPC throughput: config 1 (9121-320, 3172-1), 1 server

Run ID **VM Appl. Servers** **AIX Client Threads**	**T3112** **1** **1**	**T3122** **1** **2**	**T3132** **1** **3**	**T3142** **1** **4**
Rate (RPCs/sec)	25.0	49.3	70.7	77.3
Avg Response Time	39.9	40.4	42.3	51.6
Utilization				
Processor	35.3	65.4	92.6	99.4
Server	23.0	44.7	64.0	69.3
CPU time per RPC				
Total	14.1	13.3	13.1	12.9
Server				
Total	9.2	9.1	9.1	9.0
Virtual	8.4	8.2	8.3	8.2
CP	0.8	0.9	0.8	0.8
TCP/IP				
Total	3.9	3.4	3.4	3.3
Virtual	1.6	1.4	1.4	1.3
CP	2.3	2.0	2.0	2.0
Other	1.0	0.8	0.6	0.6
Note: All times are in milliseconds. All CPU results are from VMPRF.				

Throughput reaches a limit at about 77 RPCs per second (Table 37) because 9121-320 processor capacity has been reached. At 4 threads, processor utilization has reached 99.4%.

Average response time increases slowly as RPC throughput is increased until the processor approaches saturation at about 3 threads. The large increase when going from 3 to 4 threads reflects the fact that throughput can only increase slightly so the load applied by the fourth thread mostly causes the RPCs to experience longer delays.

CPU time per RPC decreases as the RPC throughput rate increases. This is because the constant amount of overhead from the timer-driven functions in TCP/IP, the DCECORE virtual machine, the application server, CP monitor, and other CP functions is being pro-rated across a larger number of RPCs.

The 1 thread measurement is essentially equivalent to the single thread measurement for the 0KB, protection level none RPC case shown in Table 34. Note that the total CPU times per RPC are very close (14.1 and 14.2 msec, respectively).

These throughput results are for the case of 1KB RPCs with protection level none. Since the throughput of this configuration is limited by 9121-320 processor capacity, you can estimate maximum throughput for any of the other RPC cases shown in the single thread results tables by using the ratio of CPU time per RPC for the 1KB protection level none case to the CPU time per RPC for the RPC case of interest. For example, maximum throughput for 32KB RPCs with protection level none can be estimated as:

77.3 RPCs/sec * (13.5/90.3) = 12 RPCs/sec

The 13.5 and 90.3 are total CPU time per RPC for the 1KB and 32KB protection level none cases, respectively. They are taken from Table 34.

9121-621 3172-3 Results (1 server)

Table 38. 1KB RPC throughput: config 2 (9121-621, 3172-3), 1 server

Run ID **VM Appl. Servers** **AIX Client Threads**	**T6111** **1** **1**	**T6121** **1** **2**	**T6141** **1** **4**	**T6161** **1** **6**
Rate (RPCs/sec)	49.8	97.7	154.7	158.0
Avg Response Time	20.0	20.3	25.7	37.9
Utilization				
Avg Processor	21.8	41.8	62.9	64.0
Server	30.3	59.3	91.8	93.8
CPU time per RPC				
Total	8.8	8.6	8.3	8.3
Server				
Total	6.1	6.1	6.1	6.1
Virtual	5.5	5.5	5.5	5.5
CP	0.6	0.6	0.6	0.6
TCP/IP				
Total	2.1	2.1	1.9	1.9
Virtual	0.8	0.8	0.7	0.7
CP	1.3	1.3	1.2	1.2
Other	0.6	0.4	0.3	0.3

Note: All times are in milliseconds. All CPU results are from VMPRF.

Throughput reaches a limit at about 158 RPCs per second (Table 38). In this case, throughput becomes limited by the fact that the server virtual machine[17] is fully utilized. At 4 threads, server utilization has reached 93.8%. Total system CPU utilization is only 64%, so there is additional capacity.

Compared to the 9121-320 results shown in Table 37, these 9121-621 results show significantly lower response times and CPU usage. The decreased CPU time per RPC is in proportion to the speed difference between the 9121-320 processor and one of the 9121-621 processors.

[17] The server virtual machine was configured with one virtual processor. VM DCE does not support client or server virtual machines configured as virtual multiprocessors.

TCP/IP and Other CPU usage are a smaller proportion of total CPU usage than they are for the 9121-320 measurements. For example, TCP/IP plus Other is 30% of all CPU usage for the highest utilization 9121-320 measurement. This drops to 27% for the highest utilization 9121-621 measurement. This is because the idling overhead from TCP/IP, the DCECORE virtual machine, and CP is being pro-rated across a larger number of RPCs.

9121-621 3172-3 Results (2 servers)

Table 39. 1KB RPC throughput: config 2 (9121-621, 3172-3), 2 servers

Run ID **VM Appl. Servers** **AIX Client Threads**	**T6211** **2** **2**	**T6221** **2** **4**	**T6231** **2** **6**	**T6241** **2** **8**
Rate (RPCs/sec)	98.2	180.5	207.2	214.9
Avg Response Time	20.3	22.2	29.0	37.2
Utilization				
Avg Processor	45.7	82.3	93.4	95.0
Avg Server	32.0	59.3	67.8	69.1
CPU time per RPC				
Total	9.3	9.1	9.0	9.0
Server				
Total	6.5	6.6	6.5	6.6
Virtual	5.9	6.0	6.0	6.0
CP	0.6	0.6	0.5	0.6
TCP/IP				
Total	2.3	2.2	2.2	2.2
Virtual	1.0	0.9	0.9	0.8
CP	1.3	1.3	1.3	1.4
Other	0.5	0.3	0.3	0.2

Note: All times are in milliseconds. All CPU results are from VMPRF.

The addition of a second server removed server utilization as a limiting factor. This allowed throughput to rise to 215 RPCs per second (Table 39), at which point throughput became limited by the CPU capacity of the 9121-621 system.

Note that CPU time per RPC is somewhat higher in the 2 server case compared to the 1 server case.

Figure 11 shows how the 9121-621 CPU usage is distributed among the DCE application servers, the TCP/IP virtual machine, and other sources (system overhead plus other virtual machines) for run T6241 (8 threads) in Table 39. Note that 24% of the CPU usage is from the TCP/IP virtual machine. It is possible for the utilization of the TCP/IP virtual machine to become the limiting factor on systems that have more processors and more DCE servers, or on systems that heavily use TCP/IP for other purposes in addition to DCE.

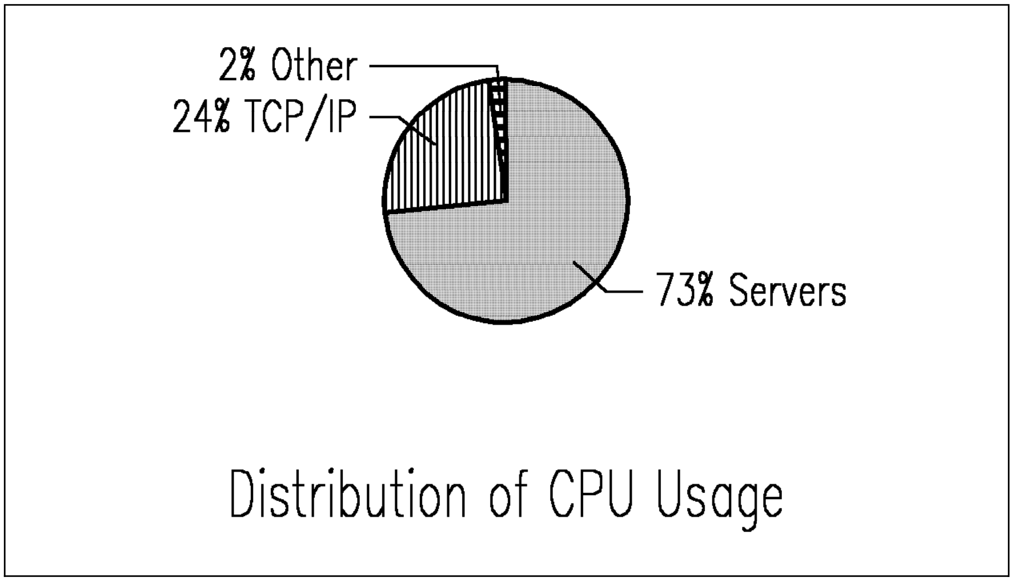

Figure 11. Distribution of 9121-621 CPU Usage (Run T6241)

9121-320 OSA-1 Results

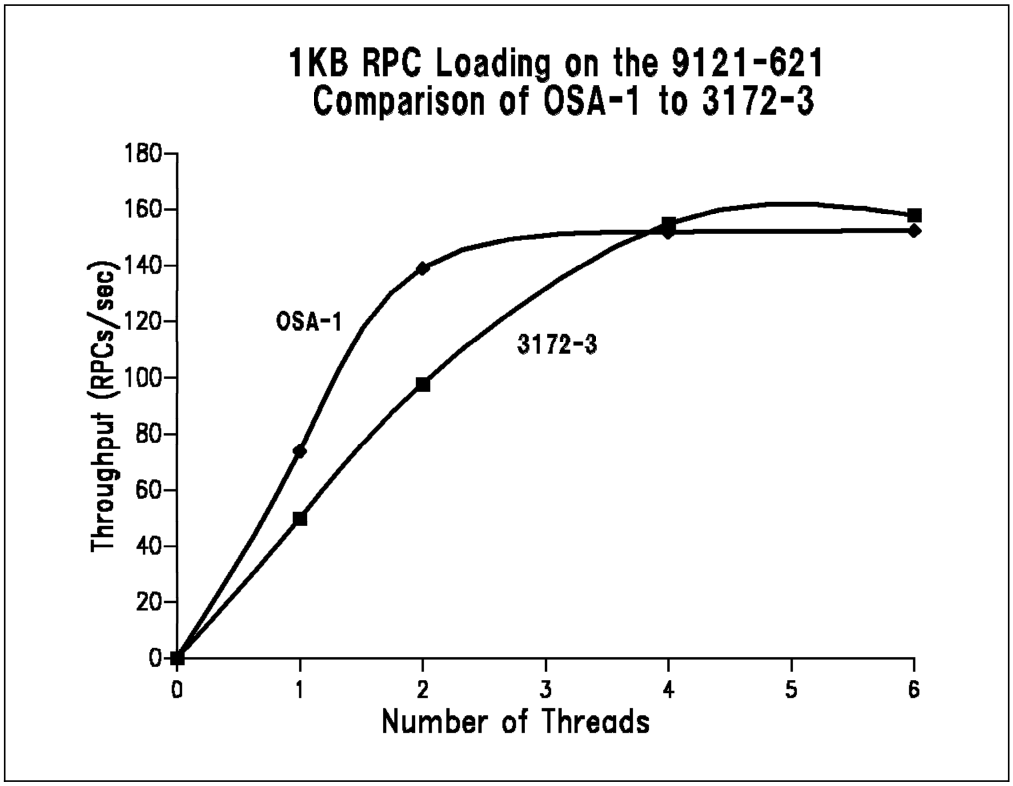

Figure 12. VM DCE throughput on a 9121-621: OSA-1/3172-3 comparison

Refer to Table 38 for the 3172-3 results and Table 40 for the OSA-1 results.

As was seen in the single thread data, the OSA-1 configuration provided better response times than the 3172-3 configuration, but at a cost of slightly more CPU usage. This determines the shape of the two curves shown in Figure 12.

At 1 and 2 threads, server CPU utilization is not yet a limiting factor. OSA-1 response times are less than the corresponding 3172-3 response times and, as a result, OSA-1 throughputs are higher. As the number of threads approaches about 4, server utilization becomes the important constraint. Because of this, response time becomes very sensitive to server CPU utilization. Because the 3172-3 configuration uses slightly less CPU time per RPC than the OSA-1 configuration, the server utilization it generates at a given RPC rate is somewhat lower. As a result, the 3172-3 curve crosses the OSA-1 curve at about 4 threads and above 4 threads the 3172-3 configuration achieves a somewhat higher throughput. This has little practical significance, however, because a properly balanced system would be operated at a much lower server utilization.

Table 40. 1KB RPC throughput: config 3 (9121-621, OSA-1), 1 server

Run ID **VM Appl. Servers** **AIX Client Threads**	**T611OSA1** **1** **1**	**T612OSA1** **1** **2**	**T614OSA1** **1** **4**	**T616OSA1** **1** **6**
Rate (RPCs/sec)	73.8	139.0	151.9	152.4
Avg Response Time	13.4	14.3	26.2	39.2
Utilization				
Avg Processor	33.1	60.8	65.1	65.4
Server	46.3	86.3	94.0	94.3
CPU time per RPC				
Total	9.0	8.8	8.6	8.6
Server				
Total	6.3	6.2	6.2	6.2
Virtual	5.6	5.6	5.6	5.6
CP	0.7	0.6	0.6	0.6
TCP/IP				
Total	2.2	2.1	2.0	2.0
Virtual	0.8	0.8	0.8	0.8
CP	1.4	1.3	1.2	1.2
Other	0.5	0.5	0.4	0.4

Note: All times are in milliseconds. All CPU results are from VMPRF.

GCS TSLICE Option

This section documents the results of a measurement made that used the new capability of altering the GCS time slice. This study was done to determine the sensitivity of this new tuning option. The VTAM virtual machine was measured using a 30 millisecond time slice and compared with the default time slice of 300 milliseconds. See "Performance Improvements" on page 9 for more discussion on this item.

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*	*PAGE*	*SPOOL*	*- Number of Volumes - 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	1900	Users	3MB/XC	100		

GCS TSLICE Option

Measurement Discussion: The following table shows that there was little effect to system performance when the GCS time slice was altered for the VTAM machine handling the CMS terminals. The key indicators of external response time (AVG LAST(T)) and internal throughput rate (ITR(H)) both show variations that are within normal run variation. The external response time decreased by 2.1% and the internal throughput improved by 0.2%. The most significant effect for this tuning parameter should be in running GCS applications that have a large number of tasks or sub-tasks.

Table 41 (Page 1 of 3). GCS time slice for the VTAM machine on a 9121-480

GCS TSLICE **Release** **Run ID**	**300** **2.1.0** **L28E190D**	**30** **2.1.0** **L28E190G**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.126	0.121	-0.005	-3.97%
NONTRIV INT	0.365	0.359	-0.006	-1.64%
TOT INT	0.286	0.279	-0.007	-2.45%
TOT INT ADJ	0.251	0.246	-0.004	-1.75%
AVG FIRST (T)	0.223	0.230	0.007	3.36%
AVG LAST (T)	0.305	0.312	0.006	2.13%
Throughput				
AVG THINK (T)	26.18	26.15	-0.02	-0.10%
ETR	58.57	58.99	0.42	0.72%
ETR (T)	66.82	66.82	0.00	0.00%
ETR RATIO	0.877	0.883	0.006	0.72%
ITR (H)	79.12	79.27	0.15	0.19%
ITR	34.69	35.01	0.32	0.92%
EMUL ITR	51.43	52.01	0.57	1.12%
ITRR (H)	1.000	1.002	0.002	0.19%
ITRR	1.000	1.009	0.009	0.92%
Proc. Usage				
PBT/CMD (H)	25.279	25.230	-0.048	-0.19%
PBT/CMD	25.293	25.293	0.000	0.00%
CP/CMD (H)	8.805	8.823	0.018	0.21%
CP/CMD	8.231	8.381	0.150	1.82%
EMUL/CMD (H)	16.474	16.407	-0.066	-0.40%
EMUL/CMD	17.061	16.912	-0.150	-0.88%
Processor Util.				
TOTAL (H)	168.90	168.58	-0.32	-0.19%
TOTAL	169.00	169.00	0.00	0.00%
UTIL/PROC (H)	84.45	84.29	-0.16	-0.19%
UTIL/PROC	84.50	84.50	0.00	0.00%
TOTAL EMUL (H)	110.07	109.63	-0.44	-0.40%
TOTAL EMUL	114.00	113.00	-1.00	-0.88%
MASTER TOTAL (H)	83.89	83.67	-0.22	-0.26%
MASTER TOTAL	84.00	84.00	0.00	0.00%
MASTER EMUL (H)	48.25	47.92	-0.33	-0.69%
MASTER EMUL	50.00	50.00	0.00	0.00%
TVR(H)	1.53	1.54	0.00	0.21%
TVR	1.48	1.50	0.01	0.88%

Table 41 (Page 2 of 3). GCS time slice for the VTAM machine on a 9121-480

GCS TSLICE **Release** **Run ID**	**300** **2.1.0** **L28E190D**	**30** **2.1.0** **L28E190G**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Storage				
NUCLEUS SIZE (V)	2756KB	2756KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	82	82	0	0.00%
PGBLPGS	54806	54767	-39	-0.07%
PGBLPGS/USER	28.8	28.8	0.0	-0.07%
FREEPGS	5657	5663	6	0.11%
FREE UTIL	0.95	0.95	0.00	-0.11%
SHRPGS	1383	1337	-46	-3.33%
Paging				
READS/SEC	638	644	6	0.94%
WRITES/SEC	439	440	1	0.23%
PAGE/CMD	16.119	16.223	0.105	0.65%
PAGE IO RATE (V)	170.800	172.500	1.700	1.00%
PAGE IO/CMD (V)	2.556	2.582	0.025	0.99%
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	8.456	8.456	0.000	0.00%
Queues				
DISPATCH LIST	36.88	36.52	-0.36	-0.97%
ELIGIBLE LIST	0.00	0.02	0.02	na
I/O				
VIO RATE	695	695	0	0.00%
VIO/CMD	10.402	10.401	0.000	0.00%
RIO RATE (V)	365	368	3	0.82%
RIO/CMD (V)	5.463	5.507	0.045	0.82%
NONPAGE RIO/CMD (V)	2.906	2.926	0.019	0.67%
DASD RESP TIME (V)	19.700	19.700	0.000	0.00%
MDC REAL SIZE (MB)	41.4	41.3	-0.1	-0.22%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	207	207	0	0.00%
MDC WRITES (I/Os)	9.58	9.52	-0.06	-0.63%
MDC AVOID	196	196	0	0.00%
MDC HIT RATIO	0.94	0.94	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.867	13.875	0.008	0.06%
DIAG/CMD	26.506	26.508	0.002	0.01%
DIAG 04/CMD	2.483	2.484	0.001	0.06%
DIAG 08/CMD	0.755	0.751	-0.004	-0.50%
DIAG 0C/CMD	1.126	1.127	0.001	0.10%
DIAG 14/CMD	0.025	0.025	0.000	-1.21%
DIAG 58/CMD	1.249	1.248	-0.001	-0.09%
DIAG 98/CMD	1.190	1.191	0.000	0.01%
DIAG A4/CMD	3.805	3.815	0.010	0.26%
DIAG A8/CMD	2.835	2.825	-0.011	-0.37%
DIAG 214/CMD	11.800	11.807	0.006	0.05%
SIE/CMD	53.100	53.099	-0.001	0.00%
SIE INTCPT/CMD	35.577	35.577	-0.001	0.00%
FREE TOTL/CMD	49.882	49.852	-0.031	-0.06%

Table 41 (Page 3 of 3). GCS time slice for the VTAM machine on a 9121-480				
GCS TSLICE **Release** **Run ID**	**300** **2.1.0** **L28E190D**	**30** **2.1.0** **L28E190G**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
VTAM Machines				
WKSET (V)	507	499	-8	-1.58%
TOT CPU/CMD (V)	4.0658	3.9992	-0.0666	-1.64%
CP CPU/CMD (V)	1.4800	1.4883	0.0083	0.56%
VIRT CPU/CMD (V)	2.5858	2.5110	-0.0748	-2.89%
DIAG 98/CMD (V)	1.191	1.191	0.000	0.02%
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Additional Evaluations

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

- “VM/ESA on the Server 500” on page 116 examines VM/ESA performance when a CMS-intensive workload is run on the PC Server 500 System/390.
- “RAMAC Array Family” on page 120 provides usage guidelines and measurement results for using the RAMAC Array Family with VM/ESA.
- “CMS Virtual Machine Modes” on page 125 compares the performance of running CMS users in 370 mode, XA mode, and XC mode virtual machines.
- “370 Accommodation” on page 131 quantifies the performance effects of running CMS users under the 370 accommodation facility.
- “Storage Constrained VSE Guest using MDC” on page 135 explores tuning considerations when using minidisk caching with VSE guests in a storage constrained environment.
- “RSCS 3.2” on page 140 compare the performance of RSCS 3.2 to RSCS 3.1.
- “DirMaint 1.5” on page 145 compares the performance of DirMaint 1.5 to DirMaint 1.4.
- “VTAM 4.2.0” on page 153 compares the performance of VTAM 4.2.0 to VTAM 3.4.1 for a CMS-intensive environment.

VM/ESA on the Server 500

The results summarized in this section are excerpted from *PC Server 500 System/390 Performance* (WSC Flash 9522), which also contains CMS plus LAN file server results, VSE/ESA CICS (native and guest) results, MVS/ESA results, and tuning guidance. The interested reader should refer to this Flash and to *IBM PC Server 500 S/390 ...Is it right for you?* for further information. See “Referenced Publications” on page 5.

Introduction

The IBM PC Server 500 System/390* is a combination of the System/390 and PC Server architectures that provides access and use of both in a single package. While the S/390* instructions execute natively on a dedicated CMOS chip on the S/390 Microprocessor Complex, the execution of the S/390′s I/O is handled by OS/2 device managers, device drivers, and S/390 channel emulation. The S/390 design point in the PC Server S/390 is unique when compared to other S/390 processors. In this implementation, S/390 devices (tapes and printers) are either channel attached (via S/370 Channel Emulator/A) or emulated on PC devices in a manner that is transparent to the S/390. In addition to emulating the S/390 I/O, the PC processor also supports OS/2 applications and advanced local area network (LAN) functions.

Measurement Description

The system configuration used for the performance measurements described in this section is listed in Table 42.

Table 42. System configuration for S/390 performance runs

Hard Drives	3
Approx GB/Drive	2.25G Fast/Wide
Average seek time (ms)	7.5
Average latency (ms)	4.17
Rotational speed (RPM)	7200
Array Stripe Width (KB)	64
Channels used on RAID adapter	1
Logical drive types	RAID-5
Logical drives per array	2
Partitions per logical drive	1
Format type	HPFS
OS/2 level	V3 Warp Fullpack GA
LAN Adapter	AutoLAN Streamer* MC 32
SCSI Adapter	F/W Streaming RAID Adapter/A

Note: Seek, latency, and RPM specifications are advertised values for individual drives in the array and were not measured here.

TPNS measurements were obtained by utilizing the AWS3172 device driver and establishing a VTAM 3172 XCA connection across an isolated Token Ring LAN.

A PC Server 500 S/390 running VM/ESA ESA Feature with VTAM was connected to a second server running VM/ESA 370 Feature with VTAM and IBM′s TPNS. TPNS simulates actual VTAM cross domain logon sessions to the target system. The simulated VTAM sessions logged onto the VSCS APPL on the target system. CMS users were then logged onto on the measured system. The users ran their scripts with TPNS measuring the end user response time through the VTAM network.

Measurements were obtained using the FS8F0R workload to get an understanding of how many CMS users can be supported in various S/390 storage sizes, while maintaining an average response time of (approximately) one second. The following conditions apply to these measurements:

- dedicated

  The PC Server 500 was dedicated to processing the CMS workload. That is, there was no other activity coming in over the LAN during the measurement period.

- HPFS write caching (lazy on)
- no RAID adapter write caching (write-through)
- VM/ESA Version 1 Release 2.2
- emulated DASD volumes:

  2 system volumes (9336)
  3 page volumes (3380)
  2 spool volumes (9336)
  3 minidisk volumes (9336)
  2 t-disk volumes (9336)

- 16KB CP trace table
- STORBUF 300 200 100
- LDUBUF 600 300 100

Measurement Discussion

The measurement results are summarized in Table 43.

Table 43. PC Server 500 S/390 Performance - CMS Workload

VM/ESA Release RUN ID	1.2.2 PC7E5075	1.2.2 PC7E5190
Environment		
S/390 Real Storage	**32MB**	**128MB**
Users	**70**	**190**
MDC BIAS	**0.1**	**0.2**
Response Time (sec)		
TRIV INT	0.50	0.30
AVG LAST (T)	0.99	1.02
Throughput		
ETR (T)	2.44	6.57
S/390 CPU Usage (msec)		
CPU/CMD (V)	154	122
CP/CMD (V)	52	36
EMUL/CMD (V)	102	86
S/390 Utilization		
TOTAL (V)	37.5	80.3
Paging		
PAGE IO RATE (V)	7.3	4.8
PAGE/CMD	18.1	5.6
PAGE IO/CMD (V)	3.0	0.7
PGBLPGS/USER	100	163
I/O		
RIO RATE (V)	23	39
MDC REAL SIZE (MB)	1.8	13.7
MDC HIT RATIO	0.69	0.86
SPM2		
PC Utilization (S)	54.5	66.5
I/O Req/sec (S)	49	61

Note: T=TPNS, V=VMPRF, S=SPM2, Unmarked=RTM

The results show that, for this workload, the number of CMS users that can be supported is mostly determined by the amount of available S/390 memory. Contention for the S/390 processor is not a significant factor until 128MB are made available, at which point S/390 processor utilization rises to 80%.

VM/ESA performs better on S/390 when steps are taken to minimize the amount of page I/O that the system has to do. Page I/Os are expensive because each I/O typically reads or writes multiple 4K pages. In addition, VM/ESA′s block paging mechanism is optimized for traditional mainframe DASD and therefore does not work as well with the device emulation used here. We reduced the page I/O rate by taking the following tuning actions:

- The minidisk cache BIAS parameter was used to reduce the amount of real storage that was used for minidisk caching. This left more real storage available to reduce paging.
- A small CP trace table was used.

We also found that the response time impact of VM/ESA page I/O tends to be reduced when emulated CKD devices are used as paging volumes and multiple page volumes are defined. This does, however, result in higher PC utilizations relative to using FBA page volumes.

The PC utilization arises from handling the S/390 I/O requests. As this utilization increases, contention for the PC processor will cause I/O service times (as seen by the S/390) to increase. The PC processor, then, is one of the resources that can limit the S/390 I/O rate that can be sustained while still providing acceptable response times.

STORBUF and LDUBUF were set to high values in order to discourage eligible list formation. This tuning action was taken during preliminary measurements at high paging rates and was found to be helpful in that environment. Non-default STORBUF and LDUBUF settings may not have been necessary for the measurements shown here because the paging rates are much lower.

Note that the OS/2 I/O rate (as reported by SPM2) is higher than the S/390 I/O rate (as reported by VMPRF). This is because the S/390 I/O emulation code will sometimes split one S/390 I/O request into multiple OS/2 I/O requests.

Lower throughputs should be expected if there is no write caching in effect (either by OS/2 HPFS or by the RAID adapter). See the Flash for further information.

RAMAC Array Family

This section provides some usage guidelines for the RAMAC Array Family and documents the results of a measurement made to observe the effects of using a RAMAC Array Subsystem for VM/ESA. For additional information see the ITSC redbook *IBM RAMAC Array Family*. For more details on using RAMAC devices in a VM/ESA environment, see *Using the IBM RAMAC Array DASD in an MVS, VM, or VSE Environment*.

The following are general performance guidelines for using RAMAC Array devices in a VM/ESA environment.

- Insure you are current on service for both hardware and software. For example, VM/ESA APAR VM59200 and VMPRF APAR VM59341 are required to process cache measurement data for the RAMAC subsystem. In addition, RTM/ESA APAR GC05363 is required for RTM to report correct I/O service times for the subsystem.
- Do proper capacity planning before migrating to a RAMAC environment. Your IBM storage specialist has modeling and sizing resources available. Use them to set appropriate performance expectations.
- Enable cache for RAMAC DASD behind 3990-3 and 3990-6 control units. Significant availability and reliability benefits are provided by use of RAID technology in RAMAC. Sufficient cache is required to get good performance when this technology is used. Make sure performance planning includes determining appropriate amount of cache and NVS (non-volatile storage). IBM storage specialists recommend enabling cache for all RAMAC DASD. This includes VM/ESA paging (more on paging below), where the traditional recommendation is to disable cache. Enabling cache is especially worthwhile when the Record Cache II feature of the 3990-6 is available.
- Placement of minidisks within a VM/ESA volume is not as important because these logical volumes are really spread across multiple real disks. Based on this, seeks data from the monitor is not as valuable as in the past, but can still be used to find heavily used minidisks.
- VM/ESA paging volumes should be low on the list of things to move to RAMAC. RAMAC is meant to increase availability and reliability; paging is temporary. Paging is higher in writes and has poor re-reference patterns. Therefore, if you have a choice between moving paging space or other data, move the other data.
- Watch for data consolidation problems. If you are merging many single density volumes onto RAMAC Arrays emulating 3390-3s, recognize that the I/O rate to that logical volume is the aggregate of the merged volumes.
- Balance I/O across drawers where possible. In addition to balancing I/O across channels, control units and devices, the RAMAC drawers are a new level. While the drawer is typically not expected to be a bottleneck, it is an additional consideration.

Workload: FS8F0R

Hardware Configuration

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

DASD: (base measurement)

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	2						2 R
3390-2	3990-3	4		2	2	16 R		

DASD: (RAMAC measurement)

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
9395-B13	9394-2	4	6	4	6			
3390-2	3990-2	4	10	2				
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. The 9394 control unit cache can not be controlled by software and is always on.

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	1900	Users	3MB/XC	100		

Measurement Discussion: A measurement was made where 16 volumes of 3390-2 DASD behind a 3990-2 control unit were moved to a 9394-2 RAMAC Subsystem configured to emulate 3380K DASD on 9395-B13 DASD. The 16 volumes consisted of 6 page, 4 spool, and 6 T-disk volumes. The FS8F CMS-intensive workload was run on both configurations. When running with the RAMAC Subsystem, the external response time (AVG LAST(T)) decreased by 5.3% and the internal throughput rate (ITR(H)) did not change significantly. For other measurement results see Table 45 on page 123.

Table 44 shows a comparison of the service time for the volumes moved to the RAMAC Subsystem. Both the spool and T-disk volumes showed significant improvement in service time, while the paging volumes showed an increase in service time. There was over a 25% reduction in the connect time component of service time, which reflects the faster data transfer capabilities of the RAMAC. This reduction in connect time was true for all three types of data. Disconnect time decreased over 40% for spool and T-disk, but increased for the paging volumes. This disconnect time change reflects the benefit of control unit cache. The 9394 caching provides much greater benefit to spool and T-disk volumes than the paging volumes.

Table 44. Volumes moved to RAMAC Array Subsystem

Volume Usage	**3990/3390 Service Time**	**RAMAC Service Time**
Spool	14.0	5.5
T-Disk	28.4	16.2
Page	19.1	21.5

Note: Service time is average for each volume usage type and is in milliseconds.

Table 45 (Page 1 of 2). RAMAC Array Subsystem on 9121-480				
RAMAC Used Release Run ID	**No VM/ESA 2.2 L27E1900**	**Yes VM/ESA 2.2 L27E1903**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.133	0.131	-0.002	-1.50%
NONTRIV INT	0.461	0.434	-0.027	-5.86%
TOT INT	0.350	0.331	-0.019	-5.43%
TOT INT ADJ	0.304	0.293	-0.011	-3.58%
AVG FIRST (T)	0.272	0.260	-0.011	-4.05%
AVG LAST (T)	0.385	0.364	-0.021	-5.33%
Throughput				
AVG THINK (T)	26.17	26.16	-0.01	-0.02%
ETR	58.16	59.23	1.07	1.84%
ETR (T)	66.94	66.87	-0.07	-0.11%
ETR RATIO	0.869	0.886	0.017	1.95%
ITR (H)	74.29	74.32	0.04	0.05%
ITR	32.99	32.95	-0.04	-0.13%
EMUL ITR	47.75	47.56	-0.19	-0.39%
ITRR (H)	1.000	1.000	0.000	0.05%
ITRR	1.000	0.999	-0.001	-0.13%
Proc. Usage				
PBT/CMD (H)	26.923	26.910	-0.013	-0.05%
PBT/CMD	26.291	26.918	0.627	2.38%
CP/CMD (H)	8.826	8.845	0.019	0.22%
CP/CMD	8.067	8.225	0.158	1.96%
EMUL/CMD (H)	18.097	18.065	-0.032	-0.18%
EMUL/CMD	18.225	18.693	0.469	2.57%
Processor Util.				
TOTAL (H)	180.23	179.94	-0.29	-0.16%
TOTAL	176.00	180.00	4.00	2.27%
UTIL/PROC (H)	90.11	89.97	-0.14	-0.16%
UTIL/PROC	88.00	90.00	2.00	2.27%
TOTAL EMUL (H)	121.15	120.80	-0.35	-0.29%
TOTAL EMUL	122.00	125.00	3.00	2.46%
MASTER TOTAL (H)	89.82	89.62	-0.20	-0.23%
MASTER TOTAL	88.00	90.00	2.00	2.27%
MASTER EMUL (H)	53.68	53.43	-0.25	-0.47%
MASTER EMUL	54.00	55.00	1.00	1.85%
TVR(H)	1.49	1.49	0.00	0.13%
TVR	1.44	1.44	0.00	-0.18%
Storage				
NUCLEUS SIZE (V)	2572KB	2576KB	4KB	0.16%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	84	84	0	0.00%
PGBLPGS	55263	55107	-156	-0.28%
PGBLPGS/USER	29.1	29.0	-0.1	-0.28%
FREEPGS	5446	5382	-64	-1.18%
FREE UTIL	0.94	0.95	0.01	1.19%
SHRPGS	1251	1323	72	5.76%

Table 45 (Page 2 of 2). RAMAC Array Subsystem on 9121-480

RAMAC Used **Release** **Run ID**	**No** **VM/ESA 2.2** **L27E1900**	**Yes** **VM/ESA 2.2** **L27E1903**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	628	633	5	0.80%
WRITES/SEC	441	445	4	0.91%
PAGE/CMD	15.969	16.121	0.152	0.95%
PAGE IO RATE (V)	180.900	181.100	0.200	0.11%
PAGE IO/CMD (V)	2.702	2.708	0.006	0.22%
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	8.634	8.928	0.294	3.40%
Queues				
DISPATCH LIST	44.41	39.65	-4.77	-10.73%
ELIGIBLE LIST	1.14	0.00	-1.14	-100.00%
I/O				
VIO RATE	648	671	23	3.55%
VIO/CMD	9.680	10.035	0.355	3.66%
RIO RATE (V)	362	366	4	1.10%
RIO/CMD (V)	5.408	5.473	0.066	1.22%
NONPAGE RIO/CMD (V)	2.705	2.765	0.060	2.21%
DASD RESP TIME (V)	19.600	19.900	0.300	1.53%
MDC REAL SIZE (MB)	40.8	40.4	-0.4	-0.87%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	185	192	7	3.78%
MDC WRITES (I/Os)	9.18	9.62	0.44	4.79%
MDC AVOID	174	181	7	4.02%
MDC HIT RATIO	0.94	0.93	-0.01	-1.06%
PRIVOPs				
PRIVOP/CMD	13.457	13.871	0.415	3.08%
DIAG/CMD	26.946	27.891	0.945	3.51%
DIAG 04/CMD	2.431	2.422	-0.009	-0.36%
DIAG 08/CMD	0.720	0.747	0.027	3.70%
DIAG 0C/CMD	1.097	1.125	0.028	2.55%
DIAG 14/CMD	0.024	0.024	0.001	3.70%
DIAG 58/CMD	1.202	1.248	0.046	3.87%
DIAG 98/CMD	1.029	1.064	0.035	3.38%
DIAG A4/CMD	3.455	3.587	0.132	3.81%
DIAG A8/CMD	2.714	2.817	0.102	3.77%
DIAG 214/CMD	13.163	13.708	0.546	4.15%
SIE/CMD	51.806	53.537	1.732	3.34%
SIE INTCPT/CMD	33.674	34.264	0.590	1.75%
FREE TOTL/CMD	48.683	49.485	0.801	1.65%
VTAM Machines				
WKSET (V)	510	508	-2	-0.39%
TOT CPU/CMD (V)	3.7594	3.7802	0.0208	0.55%
CP CPU/CMD (V)	1.4191	1.4373	0.0182	1.28%
VIRT CPU/CMD (V)	2.3403	2.3429	0.0026	0.11%
DIAG 98/CMD (V)	1.020	1.064	0.044	4.31%

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

CMS Virtual Machine Modes

This section documents the results of measurements made to observe the effects of running with virtual machine modes of 370, XA, and XC on VM/ESA 1.2.2 for a CMS-intensive workload. No exploitation of XA or XC mode was done for these measurements. With the increased virtual storage availability, there is an opportunity in many customer environments to define additional shared segments to reduce real storage requirements, thus improving system performance.

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*	*PAGE*	*SPOOL*	*- Number of Volumes - 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	1900	Users	3MB/varied	100		

Measurement Discussion: The following table shows the performance cost when migrating from a 370 mode virtual machine to an XA mode virtual machine. External response time (AVG LAST(T)) increased by 8.6%. Internal throughput (ITR(H)) decreased by 2.0%, reflecting an increase in processor usage. The majority of this increase is in the CMS interrupt handlers to process the XA mode interrupts.

Table 46 (Page 1 of 3). Migration from 370 mode to XA mode for VM/ESA 1.2.2 on the 9121-480

VM Mode **Release** **Run ID**	**370** **VM/ESA 2.2** **L27E190A**	**XA** **VM/ESA 2.2** **L27E190B**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.129	0.132	0.003	2.33%
NONTRIV INT	0.410	0.438	0.028	6.83%
TOT INT	0.317	0.335	0.018	5.68%
TOT INT ADJ	0.276	0.295	0.019	6.84%
AVG FIRST (T)	0.235	0.249	0.014	5.73%
AVG LAST (T)	0.327	0.356	0.028	8.55%
Throughput				
AVG THINK (T)	26.14	26.21	0.07	0.27%
ETR	58.19	58.86	0.67	1.15%
ETR (T)	66.76	66.79	0.03	0.05%
ETR RATIO	0.872	0.881	0.010	1.10%
ITR (H)	76.19	74.68	-1.51	-1.98%
ITR	33.22	32.92	-0.30	-0.90%
EMUL ITR	48.74	47.62	-1.12	-2.29%
ITRR (H)	1.000	0.980	-0.020	-1.98%
ITRR	1.000	0.991	-0.009	-0.90%
Proc. Usage				
PBT/CMD (H)	26.252	26.782	0.531	2.02%
PBT/CMD	26.214	26.800	0.587	2.24%
CP/CMD (H)	8.918	8.846	-0.072	-0.81%
CP/CMD	8.388	8.235	-0.154	-1.83%
EMUL/CMD (H)	17.334	17.937	0.602	3.48%
EMUL/CMD	17.825	18.566	0.740	4.15%
Processor Util.				
TOTAL (H)	175.25	178.88	3.63	2.07%
TOTAL	175.00	179.00	4.00	2.29%
UTIL/PROC (H)	87.63	89.44	1.81	2.07%
UTIL/PROC	87.50	89.50	2.00	2.29%
TOTAL EMUL (H)	115.72	119.80	4.08	3.52%
TOTAL EMUL	119.00	124.00	5.00	4.20%
MASTER TOTAL (H)	87.22	89.04	1.82	2.09%
MASTER TOTAL	87.00	89.00	2.00	2.30%
MASTER EMUL (H)	50.70	52.95	2.25	4.44%
MASTER EMUL	52.00	55.00	3.00	5.77%
TVR(H)	1.51	1.49	-0.02	-1.41%
TVR	1.47	1.44	-0.03	-1.84%

Table 46 (Page 2 of 3). Migration from 370 mode to XA mode for VM/ESA 1.2.2 on the 9121-480				
VM Mode **Release** **Run ID**	**370** **VM/ESA 2.2** **L27E190A**	**XA** **VM/ESA 2.2** **L27E190B**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Storage				
NUCLEUS SIZE (V)	2572KB	2572KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	83	84	1	1.20%
PGBLPGS	55326	55316	-10	-0.02%
PGBLPGS/USER	29.1	29.1	0.0	-0.02%
FREEPGS	5380	5388	8	0.15%
FREE UTIL	0.95	0.95	0.00	-0.15%
SHRPGS	1235	1244	9	0.73%
Paging				
READS/SEC	649	628	-21	-3.24%
WRITES/SEC	436	444	8	1.83%
PAGE/CMD	16.253	16.050	-0.202	-1.24%
PAGE IO RATE (V)	183.600	180.700	-2.900	-1.58%
PAGE IO/CMD (V)	2.750	2.706	-0.045	-1.63%
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	8.898	8.953	0.056	0.63%
Queues				
DISPATCH LIST	39.22	40.44	1.22	3.12%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	678	677	-1	-0.15%
VIO/CMD	10.156	10.136	-0.020	-0.19%
RIO RATE (V)	378	369	-9	-2.38%
RIO/CMD (V)	5.662	5.525	-0.137	-2.43%
NONPAGE RIO/CMD (V)	2.912	2.819	-0.093	-3.18%
DASD RESP TIME (V)	19.400	19.800	0.400	2.06%
MDC REAL SIZE (MB)	42.4	41.5	-0.9	-2.22%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	192	192	0	0.00%
MDC WRITES (I/Os)	9.69	9.54	-0.15	-1.55%
MDC AVOID	181	181	0	0.00%
MDC HIT RATIO	0.93	0.94	0.01	1.08%

Table 46 (Page 3 of 3). Migration from 370 mode to XA mode for VM/ESA 1.2.2 on the 9121-480

VM Mode **Release** **Run ID**	**370** **VM/ESA 2.2** **L27E190A**	**XA** **VM/ESA 2.2** **L27E190B**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
PRIVOPs				
PRIVOP/CMD	13.784	13.883	0.100	0.72%
DIAG/CMD	28.318	27.980	-0.338	-1.19%
DIAG 04/CMD	2.396	2.396	0.000	0.00%
DIAG 08/CMD	0.751	0.749	-0.001	-0.17%
DIAG 0C/CMD	1.127	1.126	-0.001	-0.05%
DIAG 14/CMD	0.025	0.024	0.000	-0.19%
DIAG 58/CMD	1.248	1.248	0.000	0.00%
DIAG 98/CMD	1.181	1.137	-0.044	-3.71%
DIAG A4/CMD	3.591	3.588	-0.003	-0.08%
DIAG A8/CMD	2.822	2.833	0.011	0.39%
DIAG 214/CMD	13.474	13.725	0.251	1.86%
SIE/CMD	54.570	53.781	-0.789	-1.45%
SIE INTCPT/CMD	36.016	34.957	-1.059	-2.94%
FREE TOTL/CMD	56.098	49.558	-6.539	-11.66%
VTAM Machines				
WKSET (V)	551	551	0	0.00%
TOT CPU/CMD (V)	3.8697	3.8179	-0.0518	-1.34%
CP CPU/CMD (V)	1.4730	1.4390	-0.0340	-2.31%
VIRT CPU/CMD (V)	2.3967	2.3789	-0.0178	-0.74%
DIAG 98/CMD (V)	1.181	1.137	-0.044	-3.75%

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

The following table shows the performance cost when migrating from a 370 mode virtual machine to an XC mode virtual machine. External response time increased by 18.6% and internal throughput decreased by 3.2%. Relative to XA mode, ITR decreased by an additional 1.2% (3.2% - 2.0%). This is mostly due to pathlength increases in the CMS interrupt handlers to save and restore access registers.

Table 47 (Page 1 of 2). Migration from 370 mode to XC mode for VM/ESA 1.2.2 on the 9121-480

VM Mode Release Run ID	370 VM/ESA 2.2 L27E190A	XC VM/ESA 2.2 L27E1909	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.129	0.130	0.001	0.78%
NONTRIV INT	0.410	0.456	0.046	11.22%
TOT INT	0.317	0.345	0.028	8.83%
TOT INT ADJ	0.276	0.307	0.031	11.13%
AVG FIRST (T)	0.235	0.272	0.037	15.50%
AVG LAST (T)	0.327	0.388	0.061	18.63%
Throughput				
AVG THINK (T)	26.14	26.15	0.01	0.04%
ETR	58.19	59.34	1.15	1.98%
ETR (T)	66.76	66.67	-0.09	-0.13%
ETR RATIO	0.872	0.890	0.018	2.11%
ITR (H)	76.19	73.76	-2.42	-3.18%
ITR	33.22	32.85	-0.37	-1.11%
EMUL ITR	48.74	47.65	-1.09	-2.23%
ITRR (H)	1.000	0.968	-0.032	-3.18%
ITRR	1.000	0.989	-0.011	-1.11%
Proc. Usage				
PBT/CMD (H)	26.252	27.113	0.862	3.28%
PBT/CMD	26.214	27.147	0.933	3.56%
CP/CMD (H)	8.918	9.009	0.091	1.02%
CP/CMD	8.388	8.399	0.011	0.13%
EMUL/CMD (H)	17.334	18.104	0.770	4.44%
EMUL/CMD	17.825	18.748	0.923	5.18%
Processor Util.				
TOTAL (H)	175.25	180.77	5.52	3.15%
TOTAL	175.00	181.00	6.00	3.43%
UTIL/PROC (H)	87.63	90.39	2.76	3.15%
UTIL/PROC	87.50	90.50	3.00	3.43%
TOTAL EMUL (H)	115.72	120.71	4.99	4.31%
TOTAL EMUL	119.00	125.00	6.00	5.04%
MASTER TOTAL (H)	87.22	89.96	2.75	3.15%
MASTER TOTAL	87.00	90.00	3.00	3.45%
MASTER EMUL (H)	50.70	53.35	2.65	5.22%
MASTER EMUL	52.00	55.00	3.00	5.77%
TVR(H)	1.51	1.50	-0.02	-1.11%
TVR	1.47	1.45	-0.02	-1.54%
Storage				
NUCLEUS SIZE (V)	2572KB	2572KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	83	84	1	1.20%
PGBLPGS	55326	55084	-242	-0.44%
PGBLPGS/USER	29.1	29.0	-0.1	-0.44%
FREEPGS	5380	5406	26	0.48%
FREE UTIL	0.95	0.95	0.00	-0.48%
SHRPGS	1235	1313	78	6.32%

Table 47 (Page 2 of 2). Migration from 370 mode to XC mode for VM/ESA 1.2.2 on the 9121-480

VM Mode Release Run ID	370 VM/ESA 2.2 L27E190A	XC VM/ESA 2.2 L27E1909	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	649	606	-43	-6.63%
WRITES/SEC	436	445	9	2.06%
PAGE/CMD	16.253	15.763	-0.489	-3.01%
PAGE IO RATE (V)	183.600	181.700	-1.900	-1.03%
PAGE IO/CMD (V)	2.750	2.725	-0.025	-0.91%
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	8.898	8.954	0.056	0.63%
Queues				
DISPATCH LIST	39.22	41.98	2.76	7.03%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	678	671	-7	-1.03%
VIO/CMD	10.156	10.064	-0.092	-0.91%
RIO RATE (V)	378	393	15	3.97%
RIO/CMD (V)	5.662	5.894	0.232	4.10%
NONPAGE RIO/CMD (V)	2.912	3.169	0.257	8.83%
DASD RESP TIME (V)	19.400	19.200	-0.200	-1.03%
MDC REAL SIZE (MB)	42.4	41.4	-1.1	-2.49%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	192	191	-1	-0.52%
MDC WRITES (I/Os)	9.69	9.55	-0.14	-1.44%
MDC AVOID	181	180	-1	-0.55%
MDC HIT RATIO	0.93	0.94	0.01	1.08%
PRIVOPs				
PRIVOP/CMD	13.784	13.906	0.122	0.89%
DIAG/CMD	28.318	27.861	-0.457	-1.61%
DIAG 04/CMD	2.396	2.401	0.005	0.19%
DIAG 08/CMD	0.751	0.752	0.001	0.16%
DIAG 0C/CMD	1.127	1.126	-0.001	-0.06%
DIAG 14/CMD	0.025	0.025	0.000	0.49%
DIAG 58/CMD	1.248	1.248	0.000	0.02%
DIAG 98/CMD	1.181	1.081	-0.100	-8.47%
DIAG A4/CMD	3.591	3.590	-0.001	-0.04%
DIAG A8/CMD	2.822	2.823	0.001	0.02%
DIAG 214/CMD	13.474	13.665	0.191	1.42%
SIE/CMD	54.570	54.070	-0.500	-0.92%
SIE INTCPT/CMD	36.016	34.605	-1.411	-3.92%
FREE TOTL/CMD	56.098	49.495	-6.603	-11.77%
VTAM Machines				
WKSET (V)	551	559	8	1.45%
TOT CPU/CMD (V)	3.8697	3.7996	-0.0701	-1.81%
CP CPU/CMD (V)	1.4730	1.4415	-0.0315	-2.14%
VIRT CPU/CMD (V)	2.3967	2.3581	-0.0386	-1.61%
DIAG 98/CMD (V)	1.181	1.081	-0.101	-8.54%

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

370 Accommodation

This section documents the results of a measurement made to observe the effects of running with 370 accommodation mode set on. A measurement was made where CMS370AC was set on for all user virtual machines. This also causes 370ACCOM to be implicitly set on for CP. This measurement was then compared to a measurement without CMS370AC set on.

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*	*PAGE*	*SPOOL*	*- Number of Volumes - 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	1900	Users	3MB/XC	100		

Measurement Discussion: The following table shows that there is little effect to system performance for this workload when 370 accommodation mode is turned on and not actually required. The key indicators of external response time (AVG LAST(T)) and internal throughput rate (ITR(H)) both show variations that are within normal run variation. The external response time improved by 1.3% and the internal throughput decreased by 0.2%.

The most significant effect is the addition of a Diagnose code X′268′ when CMS is processing interrupts. This suggests that there could be a measureable effect on system performance for workloads that have high interrupt rates. It is worth noting that when only the CP portion of the 370 accommodation facility is used (that is, 370ACCOM ON but CMS370AC OFF), these diagnose calls do not occur and performance is therefore not sensitive to the interrupt rate.

There is a performance cost associated with changing CMS virtual machines from 370 mode to XA or XC mode. Measurements on VM/ESA 1.2.2 using this same CMS-intensive workload showed a 2.0% ITR decrease when all user virtual machines were switched from 370 mode to XA mode and a 3.2% ITR decrease when the user machines were switched from 370 mode to XC mode. Refer to "CMS Virtual Machine Modes" on page 125 for additional information.

Table 48 (Page 1 of 3). 370 accommodation mode comparison on 9121-480

CMS370AC **Release** **Run ID**	**OFF** **2.1.0** **L28E190D**	**ON** **2.1.0** **L28E190H**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.126	0.125	-0.001	-0.79%
NONTRIV INT	0.365	0.363	-0.002	-0.55%
TOT INT	0.286	0.284	-0.002	-0.70%
TOT INT ADJ	0.251	0.249	-0.002	-0.81%
AVG FIRST (T)	0.223	0.218	-0.005	-2.24%
AVG LAST (T)	0.305	0.302	-0.004	-1.31%
Throughput				
AVG THINK (T)	26.18	26.17	0.00	0.00%
ETR	58.57	58.52	-0.05	-0.09%
ETR (T)	66.82	66.83	0.02	0.02%
ETR RATIO	0.877	0.876	-0.001	-0.11%
ITR (H)	79.12	78.99	-0.13	-0.16%
ITR	34.69	34.59	-0.10	-0.29%
EMUL ITR	51.43	51.57	0.14	0.27%
ITRR (H)	1.000	0.998	-0.002	-0.16%
ITRR	1.000	0.997	-0.003	-0.29%
Proc. Usage				
PBT/CMD (H)	25.279	25.319	0.041	0.16%
PBT/CMD	25.293	25.287	-0.006	-0.02%
CP/CMD (H)	8.805	8.898	0.093	1.06%
CP/CMD	8.231	8.379	0.148	1.79%
EMUL/CMD (H)	16.474	16.421	-0.052	-0.32%
EMUL/CMD	17.061	16.908	-0.154	-0.90%

Table 48 (Page 2 of 3). 370 accommodation mode comparison on 9121-480

CMS370AC **Release** **Run ID**	**OFF** **2.1.0** **L28E190D**	**ON** **2.1.0** **L28E190H**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Processor Util.				
TOTAL (H)	168.90	169.22	0.31	0.18%
TOTAL	169.00	169.00	0.00	0.00%
UTIL/PROC (H)	84.45	84.61	0.16	0.18%
UTIL/PROC	84.50	84.50	0.00	0.00%
TOTAL EMUL (H)	110.07	109.75	-0.32	-0.29%
TOTAL EMUL	114.00	113.00	-1.00	-0.88%
MASTER TOTAL (H)	83.89	83.94	0.05	0.06%
MASTER TOTAL	84.00	84.00	0.00	0.00%
MASTER EMUL (H)	48.25	47.90	-0.35	-0.72%
MASTER EMUL	50.00	50.00	0.00	0.00%
TVR(H)	1.53	1.54	0.01	0.48%
TVR	1.48	1.50	0.01	0.88%
Storage				
NUCLEUS SIZE (V)	2756KB	2756KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	82	83	1	1.22%
PGBLPGS	54806	54703	-103	-0.19%
PGBLPGS/USER	28.8	28.8	-0.1	-0.19%
FREEPGS	5657	5745	88	1.56%
FREE UTIL	0.95	0.94	-0.01	-1.53%
SHRPGS	1383	1349	-34	-2.46%
Paging				
READS/SEC	638	637	-1	-0.16%
WRITES/SEC	439	443	4	0.91%
PAGE/CMD	16.119	16.160	0.041	0.26%
PAGE IO RATE (V)	170.800	171.100	0.300	0.18%
PAGE IO/CMD (V)	2.556	2.560	0.004	0.15%
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	8.456	8.454	-0.002	-0.02%
Queues				
DISPATCH LIST	36.88	38.57	1.70	4.60%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	695	701	6	0.86%
VIO/CMD	10.402	10.489	0.087	0.84%
RIO RATE (V)	365	366	1	0.27%
RIO/CMD (V)	5.463	5.476	0.014	0.25%
NONPAGE RIO/CMD (V)	2.906	2.916	0.010	0.34%
DASD RESP TIME (V)	19.700	19.800	0.100	0.51%
MDC REAL SIZE (MB)	41.4	41.2	-0.2	-0.56%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	207	207	0	0.00%
MDC WRITES (I/Os)	9.58	9.60	0.02	0.21%
MDC AVOID	196	196	0	0.00%
MDC HIT RATIO	0.94	0.94	0.00	0.00%

Table 48 (Page 3 of 3). 370 accommodation mode comparison on 9121-480

CMS370AC **Release** **Run ID**	**OFF** **2.1.0** **L28E190D**	**ON** **2.1.0** **L28E190H**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
PRIVOPs				
PRIVOP/CMD	13.867	13.939	0.072	0.52%
DIAG/CMD	26.506	27.651	1.146	4.32%
DIAG 04/CMD	2.483	2.481	-0.002	-0.08%
DIAG 08/CMD	0.755	0.755	0.000	0.05%
DIAG 0C/CMD	1.126	1.125	-0.001	-0.10%
DIAG 14/CMD	0.025	0.025	0.000	-0.68%
DIAG 58/CMD	1.249	1.334	0.085	6.84%
DIAG 98/CMD	1.190	1.199	0.008	0.71%
DIAG A4/CMD	3.805	3.806	0.001	0.02%
DIAG A8/CMD	2.835	2.834	-0.001	-0.05%
DIAG 214/CMD	11.800	10.960	-0.840	-7.12%
DIAG 268/CMD	0.000	1.885	1.885	na
SIE/CMD	53.100	55.197	2.097	3.95%
SIE INTCPT/CMD	35.577	37.534	1.957	5.50%
FREE TOTL/CMD	49.882	49.916	0.033	0.07%
VTAM Machines				
WKSET (V)	507	527	20	3.94%
TOT CPU/CMD (V)	4.0658	4.0233	-0.0425	-1.05%
CP CPU/CMD (V)	1.4800	1.4963	0.0163	1.10%
VIRT CPU/CMD (V)	2.5858	2.5270	-0.0588	-2.27%
DIAG 98/CMD (V)	1.191	1.199	0.009	0.74%

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

Storage Constrained VSE Guest using MDC

This section examines the use of minidisk cache (MDC) in a storage constrained configuration.

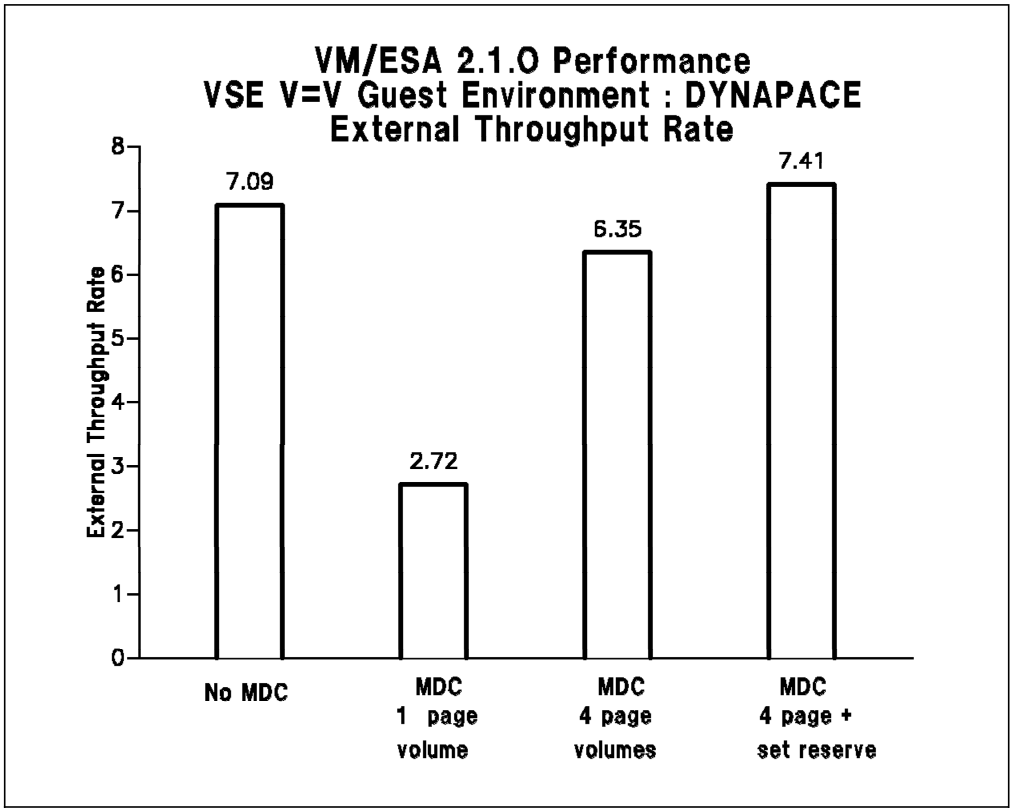

Figure 13. Storage constrained VSE guest measurements: external throughput.

In this series of measurements, 32MB of central storage was configured. The first measurement point was a V=V guest with dedicated DASD configured, so that no MDC was used. The resultant ETR was 7.09. The next measurement point used a V=V guest with MDC. In this case the VM/ESA system residence DASD had 100 cylinders of page space allocated, of which 100% was used during the measurement. A loss of about 62% in ETR occurred. By adding 3 page volumes to the system, the ETR more than doubled. This is the result of fewer I/Os due to more efficient block paging and lower I/O access times due to a lower request rate per paging device. In the final measurement point, the CP SET RESERVE command was specified for the guest (size based on the resident storage pages). SET RESERVE reduces the serial page faults in the guest machine. This resulted in an ETR increase of over 16% relative to the previous case. Overall, the use of additional page volumes and SET RESERVE resulted in a net 4.5% ETR improvement relative to the non-MDC base case.

Storage Constrained VSE Guest using MDC

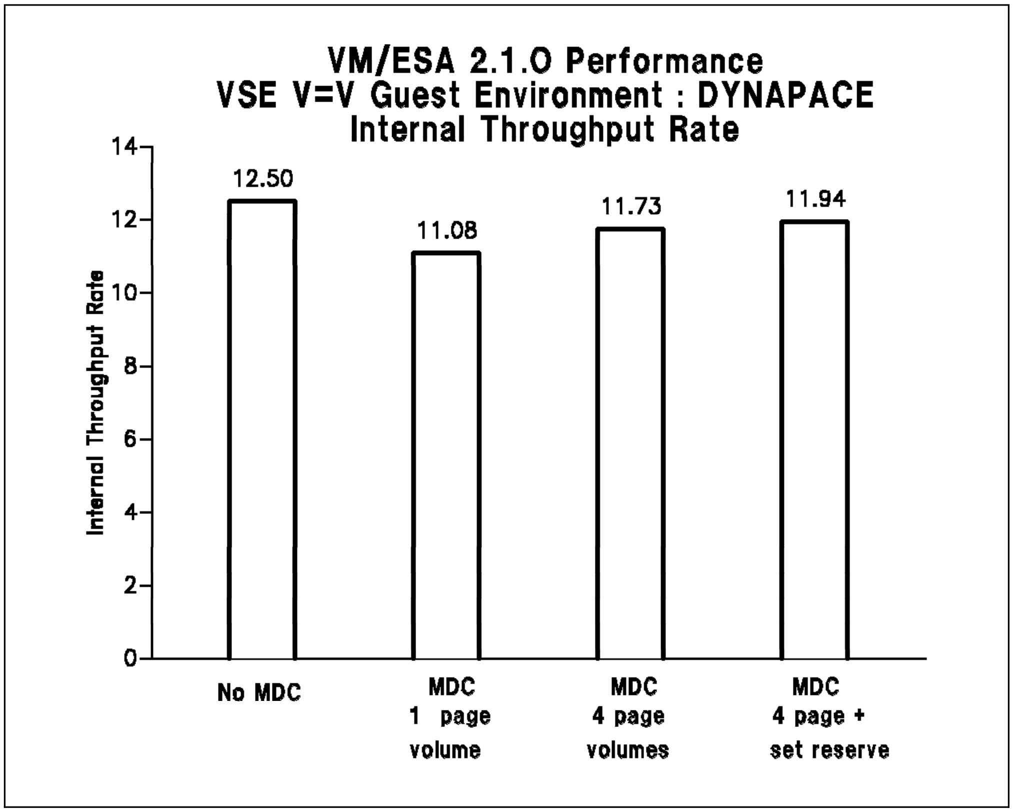

Figure 14. Storage constrained VSE guest measurements: internal throughput.

Figure 14 shows the corresponding ITRs for the four measurements described above. As a result of heavy paging, the ITR for the last measurement scenario did not exceed the measurement point where MDC was not used. Additional real storage would somewhat benefit the guest measurement where MDC is not being used, but greatly benefit the cases where MDC is used.

Workload: DYNAPACE

Hardware Configuration

Processor models: 9121-320[18]
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						1
3390-2	3990-02	4				10	2	
3380-K	3990-03	4				10		

Software Configuration

VSE version: 2.1.0

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
VSEVV	1	VSE V=V	96MB/ESA	100		IOASSIST OFF
SMART	1	RTM	16MB/370	100		
WRITER	1	CP monitor	2MB/XA	100		

For all guest measurements in this section, VSE/ESA was run in a V=V virtual machine. The VM system used for these guest measurements has a 96MB V=R area defined. For these V=V measurements, the V=R area is configured, but not used. Therefore, if the real storage configuration on the processor is 128MB, then 32MB of useable storage is available for the VM system and V=V guest. It is this effective real storage size that is shown in this section′s measurement results tables.

18 See “Hardware” on page 21 for an explanation of how this processor model was defined.

Storage Constrained VSE Guest using MDC

Table 49 (Page 1 of 2). VSE/ESA storage constrained environment on the 9121-320

MDC	**NO**	**YES**	**YES**	**YES**
Page Volumes	**1**	**1**	**4**	**4**
Set Reserve	**NO**	**NO**	**NO**	**YES**
VM/ESA Release	**2.1.0**	**2.1.0**	**2.1.0**	**2.1.0**
Run ID	**L1V88PF3**	**L1V88PF4**	**L1V88PF1**	**L1V88PF2**
Environment				
IML Mode	**ESA**	**ESA**	**ESA**	**ESA**
Real Storage	**32MB**	**32MB**	**32MB**	**32MB**
Exp. Storage	**0MB**	**0MB**	**0MB**	**0MB**
VM Mode	**ESA**	**ESA**	**ESA**	**ESA**
VM Size	**96M**	**96M**	**96M**	**96M**
Guest Setting	**V=V**	**V=V**	**V=V**	**V=V**
VSE Supervisor	**ESA**	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**	**1**
Throughput (Min)				
Elapsed Time (C)	948.0	2470.0	1057.0	907.0
ETR (C)	7.09	2.72	6.35	7.41
ITR (H)	12.50	11.08	11.73	11.94
ITR	12.44	10.88	11.77	11.95
ITRR (H)	1.000	0.886	0.939	0.955
ITRR	1.000	0.875	0.947	0.961
Proc. Usage (Sec)				
PBT/CMD (H)	4.800	5.418	5.109	5.024
PBT/CMD	4.825	5.513	5.096	5.021
CP/CMD (H)	1.248	1.864	1.589	1.508
CP/CMD	1.185	1.764	1.416	1.377
EMUL/CMD (H)	3.552	3.553	3.520	3.516
EMUL/CMD	3.640	3.749	3.681	3.644
Processor Util.				
TOTAL (H)	56.71	24.57	54.13	62.04
TOTAL	57.00	25.00	54.00	62.00
TOTAL EMUL (H)	41.96	16.11	37.30	43.42
TOTAL EMUL	43.00	17.00	39.00	45.00
TVR(H)	1.35	1.52	1.45	1.43
TVR	1.33	1.47	1.38	1.38
Storage				
NUCLEUS SIZE (V)	2776KB	2776KB	2776KB	2776KB
TRACE TABLE (V)	200KB	200KB	200KB	200KB
PGBLPGS	6163	5937	6153	6146
FREEPGS	97	100	104	107
FREE UTIL	0.59	0.57	0.58	0.59
SHRPGS	34	19	13	20
Paging				
PAGE/CMD	220.071	2844.911	2784.062	1643.938
XSTOR/CMD	0.000	0.000	0.000	0.000
FAST CLR/CMD	110.036	286.696	235.937	242.946
I/O				
VIO RATE	369.000	141.000	330.000	386.000
VIO/CMD	3123.321	3109.554	3114.375	3125.911
RIO RATE (V)	370.000	146.000	266.000	296.000
RIO/CMD (V)	3131.786	3219.821	2510.375	2397.071
DASD IO TOTAL (V)	354980	358698	286647	265495
DASD IO RATE (V)	369.77	145.81	265.41	294.99
DASD IO/CMD (V)	3129.85	3215.68	2504.84	2388.93
MDC REAL SIZE (MB)	2.2	16.3	16.0	16.3
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	0.0
MDC READS (I/Os)	0.03	86	201	235
MDC WRITES (I/Os)	0.01	42	98	115
MDC AVOID	0.00	48	119	134
MDC HIT RATIO	0.00	0.52	0.55	0.53

Table 49 (Page 2 of 2). VSE/ESA storage constrained environment on the 9121-320

MDC	**NO**	**YES**	**YES**	**YES**
Page Volumes	**1**	**1**	**4**	**4**
Set Reserve	**NO**	**NO**	**NO**	**YES**
VM/ESA Release	**2.1.0**	**2.1.0**	**2.1.0**	**2.1.0**
Run ID	**L1V88PF3**	**L1V88PF4**	**L1V88PF1**	**L1V88PF2**
Environment				
IML Mode	**ESA**	**ESA**	**ESA**	**ESA**
Real Storage	**32MB**	**32MB**	**32MB**	**32MB**
Exp. Storage	**0MB**	**0MB**	**0MB**	**0MB**
VM Mode	**ESA**	**ESA**	**ESA**	**ESA**
VM Size	**96M**	**96M**	**96M**	**96M**
Guest Setting	**V=V**	**V=V**	**V=V**	**V=V**
VSE Supervisor	**ESA**	**ESA**	**ESA**	**ESA**
Processors	**1**	**1**	**1**	**1**
PRIVOPs				
PRIVOP/CMD (R)	3125.002	3126.460	3122.455	3123.384
DIAG/CMD (R)	652.286	1212.099	686.429	627.402
SIE/CMD	13915.286	15503.661	14458.250	14171.875
SIE INTCPT/CMD	11688.840	13023.075	11855.765	11904.375
FREE TOTL/CMD	3758.143	6990.982	6134.375	5976.482

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

RSCS 3.2

This section describes the performance test activity for RSCS Version 3 Release 2. Measurement data were collected for RSCS Version 3 Release 2 and compared to RSCS Version 3 Release 1. These products are referred to as RSCS 3.2 and RSCS 3.1 throughout this section.

RSCS 3.2 supports 31 bit addressing which allows storage to be used above the 16M line. This will alleviate storage problems which required some customers to have multiple RSCS machines defined to their networks. RSCS 3.2 also has TCP/IP support for four new line drivers.

RSCS Methodology

The RSCS throughput measurements involved transferring a fixed amount of data between two systems and measuring the elapsed time and processor usage required to complete the transfer. Measurements were made while both sending and receiving data simultaneously.

The number of files transferred was held constant for all of the configurations. For each of the line drivers, 50 copies of 9 different files were sent, for a total of 450 files (see the ″RSCS Workload″ section for more details). This amount of work took about 14 minutes to complete.

Multiple measurements were made to ensure repeatability. Three measurements were done for each type of data transfer within each configuration and the results were averaged. This was done to improve the accuracy of the CPU data reported by the INDICATE USER command.

Execs were written to handle the simultaneous sending and receiving of files between both systems. The system that originally sends the files and later receives them is designated as System A and the other system that receives and resends as System B. Only System A was measured.

An exec was executed from a CMS user on System A to generate the correct number of spool files, in the desired order, and to send them to the remote user. Before sending the files, the RSCS HOLD <linkid> command was issued to ensure that no transmission activity took place before all of the files had been received by RSCS. The SENDFILE command was used to send the files. After all the files had been received by RSCS, the RSCS FREE <linkid> command was issued to initiate the transmission and begin the measurement. At that time, the exec entered into a loop to issue the 50 SMSG commands, one every 4 seconds. On System B, the user that was to receive the files was running an exec that used the WAKEUP MODULE when a RDR file was received. When a file arrived, it was received to the user′s A-disk and then sent back to RSCS and, from there, to the original user on System A who sent the files. The measurement completed when all the files had been received by the original user.

RSCS Workload

The RSCS throughput workload was built from multiple copies of a set of nine files and one SMSG command. The sizes of the files were 50 records, 200 records, and 1500 records. Each record was 80 characters in length. The compressibility of the files was also varied. RSCS compression of files is an important factor in the transmission time on TP links. In general, a record is compressible by RSCS if it contains strings of duplicate characters. The nine files are described below in the order that they are initially sent:

```
1500 records, noncompress
 200 records, semicompress
1500 records, compress
  50 records, semicompress
1500 records, semicompress
  50 records, compress
  50 records, noncompress
 200 records, compress
 200 records, noncompress
```

The records in the above files are as follows:

```
noncompress:  1234567890   (repeated 7 more times)
semicompress: 12345678901111111111   (repeated 3 more times)
compress:     11111111112222222222 ... 8888888888
```

The SMSG command was a query of the receiving RSCS virtual machine's active links and should pass immediately over the network.

```
SMSG command: SMSG <rscs1> CMD <rscs2> Q SYS A
```

SYSTEM Configuration

All of the RSCS performance measurements were made on a 9121-742 physically partitioned into two systems, each having two processors, 512MB of real storage, and 512MB of expanded storage. This resulted in a nonpaging environment for all measured cases. The measured system was connected to the other system via a 4.5 megabyte/sec 3088 (CTCA).

VM/ESA 2.1.0, VTAM 3.4.1, and TCP/IP 2.3 were used on both systems. RSCS 3.2 was at a pre-release level.

The following parameters were used for the VTAM virtual machines for both release levels of the SNANJE measurements.

```
VTAM Virtual Machine:
  - Virtual Storage Size:  64MB
  - QUICKDSP               ON
  - RESERVE                512
  - SHARE                  10000
  - Directory:             OPTION ... DIAG98
```

The following parameters were used for the RSCS virtual machines for both RSCS 3.1 and RSCS 3.2. All line drivers were run with the default number of streams (2) defined.

```
RSCS Virtual Machine:
  - Virtual Storage Size:  24MB
  - QUICKDSP               ON
```

The following parameters were used for the TCP/IP virtual machine.

```
TCP/IP Virtual Machine:
  - Virtual Storage Size:  64MB
  - QUICKDSP               ON
  - Directory:             OPTION ... DIAG98
```

RSCS 3.2 was at a pre-release level.

NJE Line Driver Results

The NJE line driver runs under the Group Control System (GCS). It provides VM with both Binary Synchronous Communication (BSC) and Channel to Channel Adapter (CTCA) line protocols to communicate with VM and non-VM NJE systems. It uses multi-streaming to transfer multiple files concurrently over the same link.

Below is the configuration used by the NJE line driver:

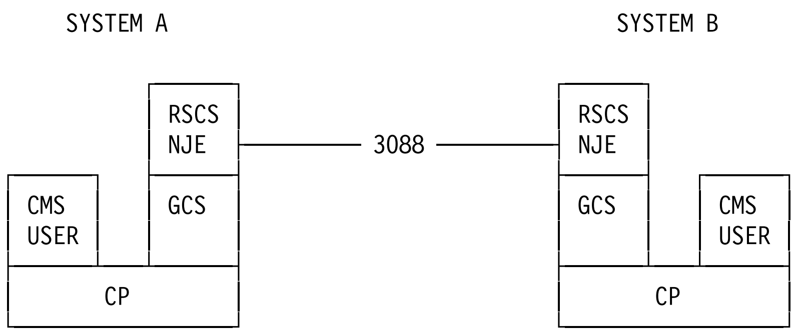

Measurement Discussion: The following table shows the throughput results when using the default buffer size of 4KB. The total and virtual CPU consumption was equivalent to RSCS 3.1. The communication I/Os (Non-spool I/Os) was within 0.4 % of RSCS 3.1.

Table 50. RSCS NJE 3088

Release	**RSCS 3.1**	**RSCS 3.2**	**Difference**	**%Difference**
TOTAL COST				
ETR	0.524	0.522	-0.002	-0.38%
Elapsed Time (sec)	858	862	4	0.47%
RSCS				
Total CPU Sec (I)	29.00	29.00	0.00	0.00%
Virtual CPU Sec (I)	10.00	10.00	0.00	0.00%
Non-Spool I/Os (I)	23255	23336	81	0.35%

Note: I=INDICATE USER

SNANJE Line Driver Results

The SNANJE line driver enables RSCS to be part of a SNA network. It runs in conjunction with GCS and VM/VTAM to provide native SNA support. It also uses multi-streaming to transfer multiple files concurrently over the same link. The Synchronous Data Link Control (SDLC) line protocol and the CTCA protocol are supported for SNANJE. Only the CTCA line protocol was tested.

Below is the configuration used by the SNANJE line driver:

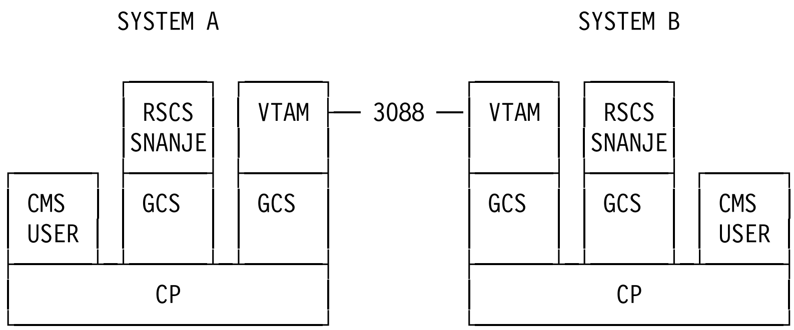

Measurement Discussion: The following table shows the throughput results when using the default buffer size of 1KB. The section of the tables called TOTAL COST is the sum of the RSCS virtual machine and the VTAM virtual machine. This allows the total cost of executing the workload to be compared. In terms of total cost, CPU usage and communication I/Os (Non-spool I/Os) were within 1% of RSCS 3.1.

Table 51. RSCS SNANJE 3088

Release	RSCS 3.1	RSCS 3.2	Difference	%Difference
TOTAL COST				
ETR	0.530	0.530	0.000	0.00%
Elapsed Time (sec)	849	849	0	0.00%
Total CPU (I)	35.70	36.00	0.30	0.84%
Virtual CPU (I)	16.67	16.63	-0.04	-0.24%
Non-Spool I/Os (I)	15146	15246	100	0.66%
RSCS				
Total CPU (I)	28.70	29.00	0.30	1.05%
Virtual CPU (I)	13.00	13.30	0.30	2.31%
Non-Spool I/Os (I)	6	8	2	33.33%
VTAM				
Total CPU (I)	7.00	7.00	0.00	0.00%
Virtual CPU (I)	3.67	3.33	-0.34	-9.26%
Non-Spool I/Os (I)	15141	15238	97	0.64%

Note: I=INDICATE USER

TCPNJE Line Driver Results

The TCPNJE line driver enables RSCS to be part of a TCP/IP network. The CTCA line protocol was also used to measure the TCP/IP environment.

Below is the configuration used by the TCPNJE line driver:

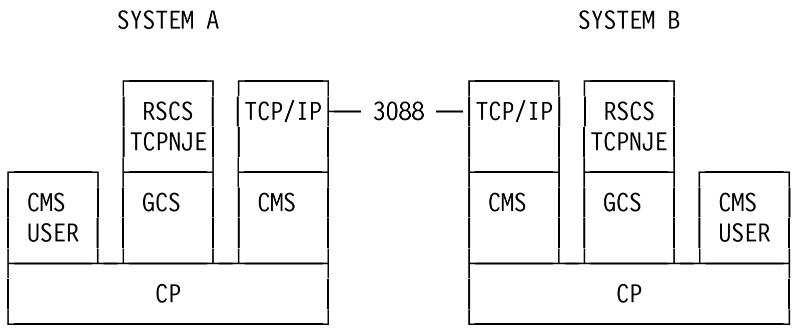

Measurement Discussion: The following table shows the throughput results when using the default buffer sizes (1KB for SNANJE and 4KB for TCPNJE). The table compares the SNANJE line driver to the TCPNJE line driver. The section of the table called TOTAL COST is the sum of RSCS and VTAM or TCP/IP. The total CPU consumption was 20.0% higher for the TCP/IP environment. This is mostly due to the TCP/IP virtual machine. The RSCS total CPU is 3.5% higher and the virtual CPU is 2.3% lower.

Table 52. RSCS TCPNJE 3088

Line Driver **Release**	**SNANJE** **RSCS 3.2**	**TCPNJE** **RSCS 3.2**	**Difference**	**%Difference**
TOTAL COST				
ETR	0.530	0.521	-0.009	-1.70%
Elapsed Time (sec)	849	864	15	1.77%
Total CPU (I)	36.00	43.25	7.25	20.01%
Virtual CPU (I)	16.63	19.50	2.87	17.26%
Non-Spool I/Os (I)	15246	19419	4173	27.37%
RSCS				
Total CPU (I)	29.00	30.00	1.00	3.45%
Virtual CPU (I)	13.30	13.00	-0.30	-2.26%
Non-Spool I/Os (I)	8	6	-2	-25.00%
VTAM or TCP/IP				
Total CPU (I)	7.00	13.25	6.25	89.29%
Virtual CPU (I)	3.33	6.50	3.17	95.20%
Non-Spool I/Os (I)	15238	19413	4175	27.40%

Note: I=INDICATE USER

DirMaint 1.5

This section documents the results of measurements made comparing DirMaint Release 1.4 to DirMaint Release 1.5. The two releases of DirMaint were run on VM/ESA 1.2.2 with the CMS-intensive workload (FS8F0R) running with the processors at approximately 80% busy. All DirMaint commands were issued from a separate virtual machine.

DirMaint has been rewritten to incorporate many outstanding customer requirements. DirMaint 1.5 has been implemented mostly in REXX with the execs being provided both in source form and compiled. Customers with the REXX run time library can use this compiled version. Two sets of measurements have been made to show the performance of both the uncompiled and compiled versions.

Workload: FS8F0R + DIR001

The workload consists of the FS8F0R CMS-intensive workload along with a DirMaint user (DIR001) and DirMaint server. The user machine issued DirMaint commands throughout the measurement period. These commands are:

```
DIRM AUTH
DIRM DROPF
DIRM HELP
DIRM REV
DIRM PW
DIRM PW
DIRM PW ?
DIRM TERM
DIRM ACCOUNT
DIRM ACCOUNT
DIRM IPL
DIRM MDISK
DIRM MDPW
```

These commands were selected because they were considered to be the more commonly used DirMaint commands. There was no study done to determine if this mix of user activity is representative of what might typically be executed. The rate of DirMaint command execution was set such that the DirMaint 1.4 server CPU utilization matched that of a local VM/ESA production system (about 0.1%).

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	PAGE	SPOOL	- Number of Volumes - 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
DIRMAINT	1	Server	16MB/XC	100		
DIR001	1	User	3MB/XC	100		
SMART	1	RTM	16MB/370	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XC	10000	560	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1650	Users	3MB/XC	100		

Measurement Discussion: The following table shows the system effects of migrating from DirMaint Release 1.4 to DirMaint Release 1.5, without the use of the compiled REXX execs. The external response time (AVG LAST(T)) increased by 4.2% and the internal throughput rate (ITR(H)) decreased by 4.1%. These deceases in the key performance indicators can be attributed to a large increase in processor usage for the DirMaint server machine and an increase in processor usage for the DirMaint user machine. The DirMaint server machine has been implemented in REXX, replacing code that had been mostly written in assembler language. High level, interpretive languages consume more resources than assembler languages.

The increase in DASD I/O RATE (V) for the DirMaint server reflects the fact that DirMaint 1.4 keeps its data in virtual storage while DirMaint 1.5 keeps its data in a file on DASD.

Table 53 (Page 1 of 3). Migration from DirMaint 1.4 on the 9121-480

DirMaint Release **VM/ESA Release** **Run ID**	**1.4** **1.2.2** **L27D1650**	**1.5** **1.2.2** **L27D1654**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1650**	**1650**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.114	0.115	0.001	0.88%
NONTRIV INT	0.322	0.342	0.020	6.21%
TOT INT	0.253	0.266	0.013	5.14%
TOT INT ADJ	0.221	0.233	0.012	5.53%
AVG FIRST (T)	0.197	0.204	0.007	3.55%
AVG LAST (T)	0.263	0.274	0.011	4.18%
Throughput				
AVG THINK (T)	26.17	26.24	0.08	0.29%
ETR	50.85	51.01	0.16	0.31%
ETR (T)	58.31	58.27	-0.03	-0.06%
ETR RATIO	0.872	0.875	0.003	0.37%
ITR (H)	73.62	70.59	-3.03	-4.11%
ITR	32.11	30.90	-1.21	-3.77%
EMUL ITR	46.64	44.24	-2.40	-5.15%
ITRR (H)	1.000	0.959	-0.041	-4.11%
ITRR	1.000	0.962	-0.038	-3.77%
Proc. Usage				
PBT/CMD (H)	27.168	28.333	1.165	4.29%
PBT/CMD	27.098	28.315	1.217	4.49%
CP/CMD (H)	9.047	9.165	0.118	1.31%
CP/CMD	8.404	8.580	0.177	2.10%
EMUL/CMD (H)	18.121	19.168	1.047	5.78%
EMUL/CMD	18.694	19.735	1.041	5.57%
Processor Util.				
TOTAL (H)	158.41	165.11	6.70	4.23%
TOTAL	158.00	165.00	7.00	4.43%
UTIL/PROC (H)	79.21	82.55	3.35	4.23%
UTIL/PROC	79.00	82.50	3.50	4.43%
TOTAL EMUL (H)	105.66	111.70	6.04	5.71%
TOTAL EMUL	109.00	115.00	6.00	5.50%
MASTER TOTAL (H)	78.49	81.83	3.34	4.25%
MASTER TOTAL	78.00	82.00	4.00	5.13%
MASTER EMUL (H)	45.84	48.71	2.88	6.27%
MASTER EMUL	47.00	50.00	3.00	6.38%
TVR(H)	1.50	1.48	-0.02	-1.41%
TVR	1.45	1.43	-0.01	-1.02%
Storage				
NUCLEUS SIZE (V)	2572KB	2572KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	84	84	0	0.00%
PGBLPGS	56071	56037	-34	-0.06%
PGBLPGS/USER	34.0	34.0	0.0	-0.06%
FREEPGS	4788	4817	29	0.61%
FREE UTIL	0.96	0.96	-0.01	-0.60%
SHRPGS	1238	1288	50	4.04%

Table 53 (Page 2 of 3). Migration from DirMaint 1.4 on the 9121-480				
DirMaint Release **VM/ESA Release** **Run ID**	**1.4** **1.2.2** **L27D1650**	**1.5** **1.2.2** **L27D1654**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1650**	**1650**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	579	583	4	0.69%
WRITES/SEC	388	393	5	1.29%
PAGE/CMD	16.584	16.749	0.164	0.99%
PAGE IO RATE (V)	156.700	159.400	2.700	1.72%
PAGE IO/CMD (V)	2.687	2.735	0.048	1.78%
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	8.901	9.061	0.160	1.79%
Queues				
DISPATCH LIST	31.74	32.98	1.25	3.92%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	609	613	4	0.66%
VIO/CMD	10.445	10.519	0.075	0.72%
RIO RATE (V)	343	349	6	1.75%
RIO/CMD (V)	5.883	5.989	0.106	1.81%
NONPAGE RIO/CMD (V)	3.195	3.254	0.059	1.83%
DASD RESP TIME (V)	19.400	19.400	0.000	0.00%
MDC REAL SIZE (MB)	42.5	41.9	-0.6	-1.33%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	169	170	1	0.59%
MDC WRITES (I/Os)	8.47	8.63	0.16	1.89%
MDC AVOID	159	160	1	0.63%
MDC HIT RATIO	0.93	0.93	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.886	13.928	0.042	0.30%
DIAG/CMD	28.609	29.413	0.804	2.81%
DIAG 04/CMD	2.726	2.730	0.004	0.16%
DIAG 08/CMD	0.755	0.770	0.015	1.98%
DIAG 0C/CMD	1.127	1.134	0.006	0.54%
DIAG 14/CMD	0.024	0.025	0.000	0.88%
DIAG 58/CMD	1.250	1.250	-0.001	-0.07%
DIAG 98/CMD	1.434	1.418	-0.017	-1.16%
DIAG A4/CMD	3.618	3.674	0.056	1.55%
DIAG A8/CMD	2.809	2.844	0.035	1.25%
DIAG 214/CMD	13.713	14.411	0.698	5.09%
SIE/CMD	55.207	55.978	0.770	1.40%
SIE INTCPT/CMD	36.437	36.385	-0.051	-0.14%
FREE TOTL/CMD	50.662	50.538	-0.125	-0.25%
VTAM Machines				
WKSET (V)	503	504	1	0.20%
TOT CPU/CMD (V)	4.0399	4.0136	-0.0263	-0.65%
CP CPU/CMD (V)	1.5340	1.5254	-0.0086	-0.56%
VIRT CPU/CMD (V)	2.5059	2.4883	-0.0176	-0.70%
DIAG 98/CMD (V)	1.435	1.419	-0.016	-1.12%

Table 53 (Page 3 of 3). Migration from DirMaint 1.4 on the 9121-480				
DirMaint Release **VM/ESA Release** **Run ID**	**1.4** **1.2.2** **L27D1650**	**1.5** **1.2.2** **L27D1654**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1650**	**1650**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
DIRMAINT Server				
CPU SECONDS (V)	2	93	91	4500.00%
CPU UTIL (V)	0.1	2.6	2.5	2500.00%
RESIDENT PAGES (V)	127	194	67	52.76%
DASD I/O RATE (V)	1.28	5.37	4.09	319.53%
DIRMAINT User				
CPU SECONDS (V)	3	30	27	900.00%
CPU UTIL (V)	0.1	0.8	0.7	700.00%
RESIDENT PAGES (V)	43	52	9	20.93%
DASD I/O RATE (V)	0.84	0.99	0.15	17.86%
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

The following table shows the system effects of migrating from DirMaint Release 1.4 to DirMaint Release 1.5, with the use of the compiled REXX execs. The external response time (AVG LAST(T)) increased by 4.0% and the internal throughput rate (ITR(H)) decreased by 3.2%. These increases are less than those observed in the previous table. The compiled execs reduced the CPU usage of the DirMaint server machine by 14% and reduced the CPU usage of the DirMaint user machine by 27%.

The percentage improvement relative to the uncompiled version is lower than often seen when comparing compiled REXX to the uncompiled equivalent. One contributing factor is the extensive use of DASD I/O in the DirMaint server. The processing associated with this I/O is the same for both the uncompiled and compiled versions. This serves to decrease the overall percentage improvement.

DirMaint 1.5

Table 54 (Page 1 of 3). Migration from DirMaint 1.4 (with compiled execs) on the 9121-480

DirMaint Release **VM/ESA Release** **Run ID**	**1.4** **1.2.2** **L27D1650**	**1.5** **1.2.2** **L27D1655**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1650**	**1650**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.114	0.115	0.001	0.88%
NONTRIV INT	0.322	0.344	0.022	6.83%
TOT INT	0.253	0.267	0.014	5.53%
TOT INT ADJ	0.221	0.233	0.013	5.76%
AVG FIRST (T)	0.197	0.203	0.006	3.05%
AVG LAST (T)	0.263	0.274	0.010	3.99%
Throughput				
AVG THINK (T)	26.17	26.19	0.02	0.06%
ETR	50.85	51.03	0.18	0.35%
ETR (T)	58.31	58.39	0.08	0.13%
ETR RATIO	0.872	0.874	0.002	0.22%
ITR (H)	73.62	71.24	-2.38	-3.23%
ITR	32.11	31.14	-0.97	-3.01%
EMUL ITR	46.64	44.68	-1.95	-4.19%
ITRR (H)	1.000	0.968	-0.032	-3.23%
ITRR	1.000	0.970	-0.030	-3.01%
Proc. Usage				
PBT/CMD (H)	27.168	28.074	0.906	3.33%
PBT/CMD	27.098	28.089	0.991	3.66%
CP/CMD (H)	9.047	9.112	0.065	0.71%
CP/CMD	8.404	8.564	0.160	1.90%
EMUL/CMD (H)	18.121	18.962	0.841	4.64%
EMUL/CMD	18.694	19.525	0.831	4.45%
Processor Util.				
TOTAL (H)	158.41	163.91	5.50	3.47%
TOTAL	158.00	164.00	6.00	3.80%
UTIL/PROC (H)	79.21	81.96	2.75	3.47%
UTIL/PROC	79.00	82.00	3.00	3.80%
TOTAL EMUL (H)	105.66	110.71	5.05	4.78%
TOTAL EMUL	109.00	114.00	5.00	4.59%
MASTER TOTAL (H)	78.49	81.16	2.67	3.40%
MASTER TOTAL	78.00	81.00	3.00	3.85%
MASTER EMUL (H)	45.84	48.24	2.40	5.24%
MASTER EMUL	47.00	50.00	3.00	6.38%
TVR(H)	1.50	1.48	-0.02	-1.25%
TVR	1.45	1.44	-0.01	-0.76%
Storage				
NUCLEUS SIZE (V)	2572KB	2572KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	84	84	0	0.00%
PGBLPGS	56071	56068	-3	-0.01%
PGBLPGS/USER	34.0	34.0	0.0	-0.01%
FREEPGS	4788	4786	-2	-0.04%
FREE UTIL	0.96	0.96	0.00	0.04%
SHRPGS	1238	1294	56	4.52%

Table 54 (Page 2 of 3). Migration from DirMaint 1.4 (with compiled execs) on the 9121-480

DirMaint Release **VM/ESA Release** **Run ID**	**1.4** **1.2.2** **L27D1650**	**1.5** **1.2.2** **L27D1655**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1650**	**1650**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	579	582	3	0.52%
WRITES/SEC	388	392	4	1.03%
PAGE/CMD	16.584	16.682	0.098	0.59%
PAGE IO RATE (V)	156.700	157.500	0.800	0.51%
PAGE IO/CMD (V)	2.687	2.698	0.010	0.38%
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	8.901	9.009	0.108	1.21%
Queues				
DISPATCH LIST	31.74	32.05	0.32	1.00%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	609	613	4	0.66%
VIO/CMD	10.445	10.499	0.054	0.52%
RIO RATE (V)	343	347	4	1.17%
RIO/CMD (V)	5.883	5.943	0.061	1.03%
NONPAGE RIO/CMD (V)	3.195	3.246	0.051	1.58%
DASD RESP TIME (V)	19.400	19.600	0.200	1.03%
MDC REAL SIZE (MB)	42.5	42.6	0.1	0.13%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	169	171	2	1.18%
MDC WRITES (I/Os)	8.47	8.62	0.15	1.77%
MDC AVOID	159	161	2	1.26%
MDC HIT RATIO	0.93	0.93	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.886	13.880	-0.006	-0.04%
DIAG/CMD	28.609	29.282	0.674	2.35%
DIAG 04/CMD	2.726	2.740	0.014	0.51%
DIAG 08/CMD	0.755	0.767	0.012	1.58%
DIAG 0C/CMD	1.127	1.134	0.007	0.59%
DIAG 14/CMD	0.024	0.024	0.000	0.04%
DIAG 58/CMD	1.250	1.249	-0.001	-0.09%
DIAG 98/CMD	1.434	1.405	-0.029	-2.06%
DIAG A4/CMD	3.618	3.684	0.066	1.82%
DIAG A8/CMD	2.809	2.832	0.023	0.82%
DIAG 214/CMD	13.713	14.273	0.560	4.08%
SIE/CMD	55.207	55.767	0.559	1.01%
SIE INTCPT/CMD	36.437	36.248	-0.188	-0.52%
FREE TOTL/CMD	50.662	50.731	0.069	0.14%
VTAM Machines				
WKSET (V)	503	505	2	0.40%
TOT CPU/CMD (V)	4.0399	4.0059	-0.0340	-0.84%
CP CPU/CMD (V)	1.5340	1.5129	-0.0211	-1.38%
VIRT CPU/CMD (V)	2.5059	2.4930	-0.0129	-0.51%
DIAG 98/CMD (V)	1.435	1.405	-0.030	-2.07%

Table 54 (Page 3 of 3). Migration from DirMaint 1.4 (with compiled execs) on the 9121-480

DirMaint Release VM/ESA Release Run ID	1.4 1.2.2 L27D1650	1.5 1.2.2 L27D1655	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1650**	**1650**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
DIRMAINT Server				
CPU SECONDS (V)	2	80	78	3900.00%
CPU UTIL (V)	0.1	2.2	2.1	2100.00%
RESIDENT PAGES (V)	127	273	146	114.96%
DASD I/O RATE (V)	1.28	5.66	4.38	342.19%
DIRMAINT User				
CPU SECONDS (V)	3	22	19	633.33%
CPU UTIL (V)	0.1	0.6	0.5	500.00%
RESIDENT PAGES (V)	43	73	30	69.77%
DASD I/O RATE (V)	0.84	1.21	0.37	44.05%

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

VTAM 4.2.0

This section documents the results of measurements made to observe the effects of migrating from VTAM 3.4.1 to VTAM 4.2.0 on VM/ESA 2.1.0 for a CMS-intensive workload.

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*	*PAGE*	*SPOOL*	*- Number of Volumes - 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	560	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1900	Users	3MB/XC	100		

VTAM 4.2.0

Measurement Discussion: The following table shows the performance results. External response time (AVG LAST(T)) was equivalent, within measurement variability. Internal throughput (ITR(H)) improved by 0.6%. This improvement was mainly due to a 3.3% drop in CPU usage by the VTAM virtual machine.

Table 55 (Page 1 of 3). Migration from VM/VTAM 3.4.1 to VM/VTAM 4.2.0 using VM/ESA 2.1.0 on the 9121-480

VTAM Release **VM Release** **Run ID**	**3.4.1** **2.1.0** **L28E190N**	**4.2.0** **2.1.0** **L28E190P**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.124	0.124	0.000	0.00%
NONTRIV INT	0.367	0.363	-0.004	-1.09%
TOT INT	0.287	0.284	-0.003	-1.05%
TOT INT ADJ	0.251	0.249	-0.002	-0.71%
AVG FIRST (T)	0.219	0.226	0.007	3.43%
AVG LAST (T)	0.303	0.309	0.006	1.98%
Throughput				
AVG THINK (T)	26.14	26.14	0.00	-0.02%
ETR	58.54	58.40	-0.14	-0.24%
ETR (T)	66.94	66.56	-0.38	-0.57%
ETR RATIO	0.874	0.877	0.003	0.33%
ITR (H)	79.03	79.52	0.49	0.62%
ITR	34.57	34.92	0.35	1.01%
EMUL ITR	51.63	52.18	0.55	1.06%
ITRR (H)	1.000	1.006	0.006	0.62%
ITRR	1.000	1.010	0.010	1.01%
Proc. Usage				
PBT/CMD (H)	25.307	25.149	-0.157	-0.62%
PBT/CMD	25.246	25.090	-0.155	-0.62%
CP/CMD (H)	8.940	8.901	-0.039	-0.44%
CP/CMD	8.365	8.263	-0.102	-1.22%
EMUL/CMD (H)	16.366	16.249	-0.118	-0.72%
EMUL/CMD	16.880	16.827	-0.053	-0.32%
Processor Util.				
TOTAL (H)	169.41	167.39	-2.01	-1.19%
TOTAL	169.00	167.00	-2.00	-1.18%
UTIL/PROC (H)	84.70	83.70	-1.01	-1.19%
UTIL/PROC	84.50	83.50	-1.00	-1.18%
TOTAL EMUL (H)	109.56	108.15	-1.41	-1.29%
TOTAL EMUL	113.00	112.00	-1.00	-0.88%
MASTER TOTAL (H)	84.12	83.08	-1.04	-1.24%
MASTER TOTAL	84.00	83.00	-1.00	-1.19%
MASTER EMUL (H)	47.90	47.31	-0.59	-1.24%
MASTER EMUL	50.00	49.00	-1.00	-2.00%
TVR(H)	1.55	1.55	0.00	0.10%
TVR	1.50	1.49	0.00	-0.30%

Table 55 (Page 2 of 3). Migration from VM/VTAM 3.4.1 to VM/VTAM 4.2.0 using VM/ESA 2.1.0 on the 9121-480

VTAM Release **VM Release** **Run ID**	**3.4.1** **2.1.0** **L28E190N**	**4.2.0** **2.1.0** **L28E190P**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Storage				
NUCLEUS SIZE (V)	2764KB	2764KB	0KB	0.00%
TRACE TABLE (V)	400KB	400KB	0KB	0.00%
WKSET (V)	86	86	0	0.00%
PGBLPGS	55058	55102	44	0.08%
PGBLPGS/USER	29.0	29.0	0.0	0.08%
FREEPGS	5588	5550	-38	-0.68%
FREE UTIL	0.96	0.92	-0.04	-4.11%
SHRPGS	1363	1294	-69	-5.06%
Paging				
READS/SEC	667	662	-5	-0.75%
WRITES/SEC	451	448	-3	-0.67%
PAGE/CMD	16.701	16.677	-0.024	-0.14%
PAGE IO RATE (V)	187.900	186.200	-1.700	-0.90%
PAGE IO/CMD (V)	2.807	2.797	-0.009	-0.34%
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	8.515	8.489	-0.026	-0.31%
Queues				
DISPATCH LIST	36.14	38.25	2.11	5.83%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	699	694	-5	-0.72%
VIO/CMD	10.442	10.427	-0.015	-0.14%
RIO RATE (V)	388	382	-6	-1.55%
RIO/CMD (V)	5.796	5.739	-0.057	-0.98%
NONPAGE RIO/CMD (V)	2.989	2.942	-0.047	-1.59%
DASD RESP TIME (V)	19.500	19.700	0.200	1.03%
MDC REAL SIZE (MB)	39.8	39.5	-0.3	-0.68%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	207	206	-1	-0.48%
MDC WRITES (I/Os)	9.49	9.40	-0.09	-0.95%
MDC AVOID	196	195	-1	-0.51%
MDC HIT RATIO	0.94	0.94	0.00	0.00%

Table 55 (Page 3 of 3). Migration from VM/VTAM 3.4.1 to VM/VTAM 4.2.0 using VM/ESA 2.1.0 on the 9121-480

VTAM Release **VM Release** **Run ID**	**3.4.1** **2.1.0** **L28E190N**	**4.2.0** **2.1.0** **L28E190P**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1900**	**1900**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
PRIVOPs				
PRIVOP/CMD	13.914	13.922	0.008	0.06%
DIAG/CMD	26.344	26.325	-0.019	-0.07%
DIAG 04/CMD	2.475	2.484	0.008	0.34%
DIAG 08/CMD	0.752	0.751	-0.002	-0.21%
DIAG 0C/CMD	1.126	1.126	0.000	0.01%
DIAG 14/CMD	0.025	0.025	0.000	-0.54%
DIAG 58/CMD	1.249	1.248	0.000	-0.02%
DIAG 98/CMD	1.243	1.220	-0.023	-1.88%
DIAG A4/CMD	3.801	3.811	0.010	0.26%
DIAG A8/CMD	2.837	2.828	-0.010	-0.34%
DIAG 214/CMD	11.599	11.602	0.003	0.02%
SIE/CMD	53.763	53.486	-0.277	-0.52%
SIE INTCPT/CMD	35.483	35.301	-0.183	-0.52%
FREE TOTL/CMD	50.043	49.895	-0.148	-0.30%
VTAM Machines				
WKSET (V)	560	569	9	1.61%
TOT CPU/CMD (V)	4.1495	4.0148	-0.1347	-3.25%
CP CPU/CMD (V)	1.5187	1.5108	-0.0079	-0.52%
VIRT CPU/CMD (V)	2.6308	2.5040	-0.1268	-4.82%
DIAG 98/CMD (V)	1.243	1.220	-0.024	-1.90%

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

Appendix A. Workloads

The workloads that were used to evaluate VM/ESA 2.1.0 are described in this appendix.

CMS-Intensive (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 56 shows the search-order characteristics of the two environments used for measurements discussed in this document.

Table 56. 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:

- A Bactrian-distribution think time averaging 30 seconds is used. (See "Glossary of Performance Terms" on page 177 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
- IBM High Level Assembler V1R1M0
- 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 57.

Table 57. 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.

CMS-Intensive (FS8F)

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)

Workload Description

PACE is a synthetic VSE batch workload consisting of 7 unique jobs representing the commercial environment. This set of jobs is replicated sixteen times, producing the *DYNAPACE* workload. The first eight copies run in eight static partitions and another eight copies run in four dynamic classes, each configured with a maximum of two partitions.

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

The programs, data, and work space for the jobs are all maintained by VSAM on separate volumes. DYNAPACE has about a 2:1 read/write ratio.

Relationship to PACEX8

In previous VM/ESA performance reports, PACEX8 was used as the batch workload for VSE guest measurements. DYNAPACE differs from PACEX8 in the following respects:

- It runs 16 copies of the PACE jobstream (instead of 8 copies).
- The additional 8 copies are run in dynamic partitions.
- For those jobs that run in dynamic partitions, it uses VSE virtual disk in storage for the COBOL compiles and the sort work files.
- The number of elliptical calculations in the FORTRAN job is increased from 4 iterations to 19 for increased processor loading.

Measurement Methodology

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 eight copies of PACE. Four dynamic classes, each with two partition assignments, run another eight copies of PACE.

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 the workload is ready to run, the following preparatory steps are taken:

- CICS/ICCF is active but idle

- VTAM is active but idle
- VSE/EXPLORE is active
- The LST queue is emptied (PDELETE LST,ALL)
- The accounting file is deleted (J DEL)

Once performance data gathering is initiated for the system (hardware instrumentation, CP MONITOR, and RTM), the workload is started by releasing all of the batch jobs into the partitions simultaneously using the POWER command, PRELEASE RDR,*Y.

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. ETR is calculated as the total elapsed time from the moment the jobs are released until the last partition is waiting for work.

At workload completion, the ITR is calculated by dividing the number of batch jobs by average processor busy time. The processor busy time is calculated as elapsed (wall clock) time multiplied by average processor busy percent divided by 100. The ITR value is multiplied by 60 to represent jobs per CPU busy minute.

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.

Four independent CICS partitions are run to effectively utilize the measured processor. The storage configuration for this workload is 96MB central storage and no expanded storage. Each of the four CICS partitions accesses 8 VSAM KSDS files. Measurements are taken at the 70% and 90% processor utilization points.

CICS is measured by logging on a predefined number of users, each of which starts running commands from 1 of 12 possible scripts. Once the system reaches a steady state condition, the think time is adjusted to provide a transaction rate that will cause the processor to reach the target utilization level.[19] 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 2.1.0, CICS/VSE* 2.3.0, ACF/VTAM* 4.2.0. 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 additional system performance data.

The workload executes a combination of COBOL and assembler applications to produce a 40% read and 60% 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:

[19] Think time was 11 seconds for all measurements in this report.

Table 58. 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 Reservations	2	1			1				3
	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 Specification	18	8	10						10
	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

38 DASD volumes (including DOSRES and SYSWK1) are required to run this workload. Each CICS (CICS01 - CICS04) has its own set of 8 dedicated volumes for VSAM data files. There should be at least two CHPIDS to each string of data volumes.

At every measurement point, a CICSPARS report is generated for each of the four CICS workload systems. To determine the total transaction count, which is used to calculate the ITR, the TOTAL TASKS SELECTED fields from all CICSPARS reports are added together.

The ITR is calculated as

$$\frac{\text{transactions}}{\text{processor busy seconds}}$$

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:

VM/ESA 1.2.2 ..

- 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.

VM/ESA 2.1.0 ..

- VMLIB contains the CSL code.
- DMSRTSEG contains the REXX runtime library.

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 59 covers all performance measurement results except for migration results, which are covered by Table 22 on page 75. 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*

210 *VM/ESA Version 2 Release 1.0 Performance Report* (this document)

See "Referenced Publications" on page 5 for more information on these reports.

Table 59 (Page 1 of 3). Sources of VM performance measurement results

Table 59 (Page 2 of 3). 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 11 210
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
Processor-Constrained Environment	22
RAMAC Array Family	210
370 Accommodation	210
RSCS 3.2	210
DirMaint 1.5	210
VTAM 4.2.0	210
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
Using MDC with a Storage Constrained VSE Guest	210
Processor Capacity	
3090-600J	10
9121-480	11
9021-720	11
9021-900	20
9121-742	21
9021-941	22
PC Server 500	210

Table 59 (Page 3 of 3). Sources of VM performance measurement results

Subject	Report(s)
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
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 from the VSE console timestamps or from the CICSPARS reports (CICS transaction performance data).

(H). Denotes data from the internal processor instrumentation tools.

(I). Denotes data from the CP INDICATE USER command.

(Q). Denotes data from the SFS QUERY FILEPOOL STATUS command.

(QT). Denotes data from the CP QUERY TIME 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.

(S). Identifies OS/2 data from the licensed program, System Performance Monitor 2 (SPM2).

(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 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.

BFS. Byte File System

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.

CPU SECONDS (V). Total CPU time, in seconds, used by a given virtual machine. This is the Total CPU Seconds column in VMPRF′s USER_RESOURCE_UTIL report.

CPU UTIL (V). The percentage of total system CPU time that is consumed by a given virtual machine. This is the CPU Pct column in VMPRF′s USER_RESOURCE_UTIL 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 per second that are issued to DASD on behalf of a given virtual machine. This is the DASD Rate While Logged column in VMPRF′s USER_RESOURCE_UTIL report.

For PACE measurements, the number of real SSCH or RSCH instructions per second issued to DASD on behalf of the VSE guest. 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.

DCE. Distributed Computing Environment. An industry standard for implementing distributed computing.

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 268/CMD. The number of DIAGNOSE code X′268′ instructions used per command. DIAGNOSE code X′268′ is used by the CMS370AC function. This is calculated by:

$$\frac{DIAG\_268}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_268 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, 56, and 112 for the PACEX4, PACEX6, PACEX8, and DYNAPACE 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. 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:

$$\frac{1}{(1000 \times \text{ETR (T)})} \times \sum_{(f \in filepools)} \frac{Total\_BIO\_Request\_Time_f}{\text{SFSTIME}_f}$$

Total_BIO_Request_Time is from the QUERY FILEPOOL STATUS command.

IO/CMD (Q). SFS file I/Os per command. This is calculated by:

$$\frac{1}{\text{ETR (T)}} \times \sum_{f \in filepools} \frac{Total\_IO\_Requests_f}{\text{SFSTIME}_f}$$

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:

$$6000 \times \frac{Jobs}{\text{Elapsed time (H)} \times \text{UTIL/PROC (H)}}$$

Jobs is the number of jobs run in the workload. The values of *Jobs* are 28, 42, 56, and 112 for the PACEX4, PACEX6, PACEX8, and DYNAPACE 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.

I/O Req/sec (S). I/O requests per second. This is Access Rate, taken from the SPM/2 DISK report, summed over all the Physical IDs that the S/390 workload is using.

k. Multiple of 1000.

Kb. Kilobits. One kilobit is 1024 bits.

KB. Kilobytes. One kilobyte is 1024 bytes.

LUW Rollbacks (delta). The total number of SFS logical units of work that were backed out during the measurement interval, summed over all production file pools. This is calculated by:

$$\sum_{f \in \text{filepools}} LUW\_Rollbacks_f$$

LUW_Rollbacks is from the QUERY FILEPOOL STATUS command.

MASTER EMUL. Total emulation state utilization for the master processor. For uniprocessors this is the same as TOTAL EMUL and is generally not shown. 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 1.2.2, 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 1.2.2 and higher, 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 1.2.2, 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 1.2.2 and higher, 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 1.2.2.

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 1.2.2.

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 1.2.2.

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 1.2.2.

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 1.2.2.

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 1.2.2.

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 1.2.2, this is MDNE for the total RTM time interval (<-..) on the RTM SYSTEM screen, divided by 256.

For VM/ESA 1.2.2 and higher, 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 1.2.2, 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 1.2.2 to remove these restrictions. With VM/ESA 1.2.2, 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 average response time for 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.

Non-Spool I/Os (I). Non-spool I/Os done by a given virtual machine. This is calculated from INDICATE USER data obtained before and after the activity being measured. The value shown is final IO - initial IO.

NPDS. No Page Data-Set. A VSE/ESA option, when running on VM/ESA as a V=V guest, that eliminates paging by VSE/ESA for improved efficiency. All paging is done by VM/ESA.

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.

OSA. IBM S/390 Open Systems Adapter. An integrated S/390 hardware feature that provides an S/390 system with direct access to Token Ring, Ethernet, and FDDI local area networks.

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:

$$\frac{\text{READS/SEC} + \text{WRITES/SEC}}{\text{ETR (T))}}$$

For the PACE workload:

$$60 \times \frac{\text{READS/SEC} + \text{WRITES/SEC}}{\text{ETR (H)}}$$

PAGE IO RATE (V). The number of real SSCH or RSCH instructions issued on behalf of system paging.

Glossary of Performance Terms

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:

$$6000 \times \frac{\text{TOTAL (H)}}{\text{ETR (H)}}$$

PC Utilization (S). PC processor utilization. This is Processor % Util from the CPU section of the SPM2 report.

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:

$$\frac{\text{PGBLPGS}}{\text{USERS}}$$

POSIX. A set of IEEE standards that define a standard set of programming and command interfaces based on those provided by the various UNIX implementations.

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:

$$\frac{1}{(\text{ETR (T)}) \times \mathit{RTM\_Time}} \times \sum_{x \in \text{privops}} \mathit{TOTALCNT}_x$$

For the PACE workload:

$$\frac{60}{\text{ETR (H)} \times \mathit{RTM\_Time}} \times \sum_{x \in \text{privops}} \mathit{TOTALCNT}_x$$

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.

RAID. Redundant array of independent DASD.

RAMAC. A family of IBM storage products based on RAID technology. These include the RAMAC Array Subsystem and the RAMAC Array DASD.

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

RESIDENT PAGES (V). The average number of nonshared pages of central storage that are held by a given virtual machine. This is the Resid Storage Pages column in VMPRF′s USER_RESOURCE_UTIL report.

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.

RPC. Remote Procedure Call. A client request to a service provider located anywhere in the network.

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Sac\_Calls_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{Total\_Filepool\_Requests_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{1}{\sum_{f \in \text{filepools}} \frac{Checkpoints\_Taken_f}{\text{SFSTIME}_f}}$$

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.

SPM2. System Performance Monitor 2. An IBM licensed program that collects and reports performance data for an OS/2 system.

STARS. System Trace Analysis Reports. Provides various reports based on the analysis of instruction trace data.

S/390 Real Storage. On an IBM PC Server 500 system, the amount of real storage that is available to the System/390 processor.

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 CPU (I). Total CPU time, in seconds, used by a given virtual machine. This is calculated from INDICATE USER data obtained before and after the activity being measured. The value shown is final TTIME - initial TTIME.

Total CPU (QT). Total CPU time, in seconds, used by a given virtual machine. This is calculated from QUERY TIME data obtained before and after the activity being measured. The value shown is final TOTCPU - initial TOTCPU.

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.

Total Time (QT). Elapsed time, in seconds. This is calculated from QUERY TIME data obtained before and after the activity being measured. The value shown is the final CONNECT timestamp - the initial CONNECT timestamp, converted to seconds.

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 average response time 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:

$$60 \times \frac{\text{VIO RATE}}{\text{ETR (H)}}$$

Virtual CPU (I). Virtual CPU time, in seconds, used by a given virtual machine. This is calculated from INDICATE USER data obtained before and after the activity being measured. The value shown is final VTIME - initial VTIME.

Virtual CPU (QT). Virtual CPU time, in seconds, used by a given virtual machine. This is calculated from QUERY TIME data obtained before and after the activity being measured. The value shown is final VIRTCPU - initial VIRTCPU.

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.

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:

$$\frac{\text{XSTOR IN/SEC} + \text{XSTOR OUT/SEC}}{\text{ETR (T)}}$$

For the PACE workload:

$$60 \times \frac{\text{XSTOR IN/SEC} + \text{XSTOR OUT/SEC}}{\text{ETR (H)}}$$