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

Document Number GC24-5801-01

Nov 18, 1997 (first printed Dec 11, 1996)

VM Performance Group

IBM Corporation
Endicott, New York
ernsberw@vnet.ibm.com

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 2.2.0 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 2.2.0. See the IBM Programming Announcement for VM/ESA 2.2.0 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
DirMaint
DisplayWrite
ECKD

Enterprise System/9000
ES/9000
IBM
MVS/ESA
OfficeVision
OpenEdition
OS/2
PR/SM
Processor Resource/Systems Manager
RACF
RAMAC
RS/6000
SAA
System/390
S/390
Virtual Machine/Enterprise Systems Architecture
VisualGen
VM/ESA
VM/XA
VSE/ESA
VTAM
3090

Acknowledgements

The following people contributed to this report:

Bill Bitner
Marty Horan
Gary Eheman
Wes Ernsberger
Bill Guzior
Chuck Morse
Bert Paul

Abstract

The *VM/ESA Version 2 Release 2.0 Performance Report* summarizes the performance evaluation of VM/ESA 2.2.0. Measurements were obtained for the CMS-intensive, VSE guest, and OfficeVision environments on various ES/9000 processors. Discussion covers the performance changes in VM/ESA 2.2.0, the performance effects of migrating from VM/ESA 2.1.0 to VM/ESA 2.2.0, and additional evaluations.

Note: This report was first made available 12/96. It has since been updated to include a section on the SFS A-directory support available with OfficeVision 1.4.0.

Referenced Publications

The following publications and documents are referred to in this report.

- *VM/ESA: Performance,* SC24-5782
- *VM/ESA: CMS File Pool Planning, Administration, and Operation*, SC24-5751
- *VM/ESA Performance on the P/390 and R/390*, http://www.vm.ibm.com/perf/

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
- *VM/ESA Version 2 Release 1.0 Performance Report,* GC24-5801

Much additional VM/ESA performance information is available on the VM/ESA Performance page at http://www.vm.ibm.com/perf/.

[1] This report is no longer orderable. LIST38PP softcopy is available as VM10PERF PACKAGE on VMTOOLS. Or send a note to ernsberw@vnet.ibm.com requesting a copy.

Summary of Key Findings

This report summarizes the performance evaluation of VM/ESA* Version 2 Release 2.0. Measurements were obtained for the CMS-intensive, OfficeVision*, and VSE guest 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 2.2.0 includes a number of performance enhancements (page 3). Some changes have the potential to adversely affect performance (page 5). Lastly, a number of changes were made that affect VM/ESA performance management (page 6).

Migration from VM/ESA 2.1.0: Benchmark measurements show the following performance results for VM/ESA 2.2.0 relative to VM/ESA 2.1.0:

CMS-intensive Internal throughput rate (ITR) improvements ranging from 0.8% to 2.2% and response time improvements from 2% to 17% were observed for the four measured CMS-intensive environments (page 11). The largest improvements occurred in environments with expanded storage used for paging. These improvements are primarily due to:

- Reduced CMS Working Set Size

  CMS working set size decreased by 3 to 7 pages for the measured environments, while total pages per user decreased by 24 to 30 pages as a result of this improvement.
- EXEC 2 to Compiled REXX Conversion
- GCS Preallocated Storage

See “Performance Improvements” on page 3 for more information on these and other enhancements.

OfficeVision ITR and response time are equivalent for the measured IBM Office Benchmark (IOB) environment (page 32).

VSE guest Measurements showed equivalent ITR and elapsed times for the DYNAPACE I/O-intensive batch workload (page 36). Both V=R and V=V guests were measured.

Migration from Other VM Releases: The performance measurement data in this report can be used in conjunction with similar data in the five 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 2.2.0 (page 40).

Additional Evaluations

VMCF’s master processor dependency has been removed in VM/ESA 2.2.0. Measurement results demonstrate that this can result in increased throughput for high n-way environments that make significant use of VMCF and are currently constrained on the master processor (page 48).

Measurement results show an ITR ratio of 0.95 for the 9672-R53 relative to the 9121-742 for the FS8F0R CMS-intensive workload (page 49). This ITRR is somewhat lower than those measured for the PD4 and HT5 workloads used for

VM/ESA in the Large Systems Performance Reference (LSPR). Instead, the observed ITRR for FS8F0R is similar to the average ITRR of 0.93 reported for the various MVS/ESA* LSPR workloads. The MVS/ESA LSPR data should be used in addition to the VM/ESA LSPR data when sizing a migration from a 3090, 9021, or 9121 processor to a 9672 or 2003 processor.

A set of CMS-intensive measurements was obtained on the PC Server 520 and the RS/6000 591 (page 53). Relative to earlier PC Server 500 data when running the same number of CMS users (190), the 520 results showed a 28% response time improvement, while the 591 results showed a 50% response time improvement. However, the results indicate that the 520 and 591 support only slightly more users than the 500 because capacity is gated by the S/390* Microprocessor Complex, which is the same on all three systems.

A CMS-intensive measurement on a 9672-R53 configured with 1GB of central storage and 1GB of expanded storage showed overall performance that was similar to an equivalent measurement with all 2GB configured as central storage (page 54). A number of second order performance differences were observed.

RPC throughput results show equivalent performance when migrating from TCP/IP 2.3.0 to TCP/IP 2.4.0 (page 58). The measured workload exercises the mainline path through UDP and IP in the TCP/IP protocol stack.

IOB workload results show equivalent performance when migrating from OfficeVision 1.2.0 to OfficeVision 1.3.0 (page 59).

OfficeVision 1.4.0 supports the use of an SFS directory as filemode A. IOB measurement results are provided for the case where an SFS directory is used instead of a minidisk for filemode A. In addition, a method is provided for estimating the percentage increase in processor usage for moving a given amount of minidisk activity to SFS filecontrol directories (page 63).

Migration of CSP applications to VisualGen* can result in large processor usage reductions because the application code produced by VisualGen is compiled. The two measured test applications, Heavy Math and List Inquiry, showed processor usage reductions of 73% and 55% respectively (page 69).

Changes That Affect Performance

This chapter contains descriptions of various changes to VM/ESA 2.2.0 that affect performance. This information is equivalent to the information on VM/ESA 2.2.0 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." However, some have the potential to adversely affect performance. These are listed under "Performance Considerations" on page 5. 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.

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 6 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
  - Removal of VMCF Master Processor Dependency
- CMS
  - EXEC 2 to Compiled REXX Conversion
  - PEEK of Disk Dump Files
  - FILELIST using the FILELIST Option
  - CMSINST above the 16M Line
  - Reduced CMS Working Set Size
- Other
  - GCS Preallocated Storage
  - CMS GUI Improvements

Removal of VMCF Master Processor Dependency

The VMCF modules have been changed to have full MP locking capability. As a result, VMCF functions no longer have to run on the master processor. This can result in improved throughput and response times for systems that are currently master processor constrained and that run applications that use VMCF (such as TCP/IP and OfficeVision/VM). In extreme cases of high VMCF usage and high master processor constraint, this improvement can also result in improved processor usage efficiency. See "Removal of VMCF Master Processor Dependency" on page 48 for measurement results.

EXEC 2 to Compiled REXX Conversion

RDRLIST, FILELIST, PEEK, and their associated Xedit macros have been rewritten in REXX, performance-optimized, and compiled. This significantly reduces the amount of processor time used by these functions. Processor usage reductions ranging from 22% to 27% have been observed for RDRLIST and PEEK. The percentage reduction in processor usage for FILELIST is inversely proportional to the number of files being displayed. For example, a 17% reduction was observed for the case of 10 files, while a 9% reduction was observed for 1000 files.

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 into CMSINST so as to retain the performance advantages.

PEEK of Disk Dump Files

The elapsed time and processor time required to use the PEEK exec to view spool files created by the DISK DUMP command have been greatly reduced. The larger the file, the larger the percentage improvement. Percentage reductions exceeding 90% have been observed for disk dump files having more than 8000 records.

FILELIST using the FILELIST Option

The implementation of the FILELIST option has been changed significantly, resulting in improved performance. Decreases in processor usage ranging from about 10% (few files) to about 40% (many files) have been observed.

CMSINST above the 16M Line

The CMSINST logical shared segment has been moved above the 16 MB line, thus making more virtual storage available for data and applications that must reside below the line.

The following additional CMS system files have been moved into CMSINST:

ALIALIST EXEC	PROFALIA XEDIT	SAUTH XEDIT
AUTHLIST EXEC	PROFAUTH XEDIT	SDIR XEDIT
CSLLIST EXEC	PROFCATL XEDIT	SMODE XEDIT
CSLMAP EXEC	PROFCLST XEDIT	VMLCAT XEDIT
DIRLIST EXEC	PROFCSLM XEDIT	X$CLST$X XEDIT
MACLIST EXEC	PROFDLST XEDIT	X$CSLM$X XEDIT
NAMES EXEC	PROFFDAT XEDIT	X$DLST$X XEDIT
NOTE EXEC	PROFFSEA XEDIT	X$MLST$X XEDIT
SENDFILE EXEC	PROFFSHR XEDIT	X$NAME$X XEDIT
VMLINK EXEC	PROFMLST XEDIT	X$NVML$X XEDIT
X$CATL$X EXEC	PROFNOTE XEDIT	X$ONAM$X XEDIT
X$NDIR$X EXEC	PROFSEND XEDIT	X$SEND$X XEDIT
$@VML$NK XEDIT	PROFVMLK XEDIT	

When these functions are used, this change decreases per-user real storage requirements and eliminates the processor time and I/Os that were required to load them from the S-disk into memory.

The only sizable CMS areas that remain below the 16 MB line are 1 MB of the shared portion of the CMS saved system, the nonshared portion of the CMS saved system (72 KB), and file status table (FST) entries (including the HELP segment). FSTs for read-mostly shared minidisks can be removed from the

user's primary address space by replacing these minidisks with one or more SFS dircontrol directories residing in VM data spaces.

Reduced CMS Working Set Size

CMS working set size and DASD page slot usage have decreased for the case where CMS multitasking functions have not been used during the CMS session. These improvements translate into reduced paging, lower processor usage, and improved response times. See "CMS-Intensive" on page 11 and "OfficeVision" on page 32 for quantification of benefits for the measured CMS-intensive and OfficeVision environments.

The amount of improvement will typically be larger on systems that use expanded storage for paging. The removal of large numbers of one-time referenced pages has freed up much space in expanded storage that can now be used for other pages that are more likely to be re-referenced. This improved expanded storage efficiency allows for a higher percentage of total paging to be done in expanded storage, resulting in a sizable DASD page I/O reduction.

This improvement fixes a problem that was introduced in VM/ESA 2.1.0 and therefore only applies to migrations from VM/ESA 2.1.0.

GCS Preallocated Storage

Most of the OS GETMAIN/FREEMAIN requests made by GCS have been eliminated through the use of preallocated storage, resulting in decreased processor usage. This change resulted in a 0.3% decrease in total system processor usage for the FS8F CMS-intensive workload.

CMS GUI Improvements

CMS GUI performance has been improved in the following areas:

- Xedit

  The performance of Xedit when invoked from the CMSDESK GUI application has been improved substantially. Average response time decreased by over 40% and host processor CPU usage decreased by more than 60% for a set of example Xedit interactions.

- DtEventNotify

  DtEventNotify is one of the functions provided by the distributed GUI toolkit (DT) API. The amount of host and workstation processing done by DtEventNotify has been reduced considerably (60% in one example) in the version of DT that is shipped with VM/ESA 2.2.0.

The DtEventNotify improvement, and most of the Xedit improvements, first appeared in the 6/96 Recommended Service Update to VM/ESA 2.1.0.

Performance Considerations

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

Potential Shared Segment Overlaps

Some of the shared segments provided with VM/ESA have been moved to new locations. While installing VM/ESA 2.2.0., check to make sure that this has not caused any overlaps with other shared segments in your system.

Performance Management

These changes affect the performance management of VM/ESA.

- Monitor Enhancements
- Effects on 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 2.2.0.

- Scheduler Data

  The monitor previously recorded dispatch list priority, which is relative to ATOD (artificial time of day), but did not record the ATOD value. This made the priority information less useful. ATOD and ATOD2 (associated with limit list processing) have been added to records D0/R10, D2/R4, D2/R5, and D2/R6.

- Master Processor Data

  The dispatcher uses a processor local dispatch vector (PLDV) construct to manage work from the dispatch list associated with a given processor. There is a PLDV for each processor. The master processor also has a second PLDV for master-only work. The monitor uses high frequency sampling to record the number of tasks on each of the PLDVs. For the master processor, this is the sum of both PLDVs. The number of tasks on the master-only PLDV is now recorded in D5/R3 in addition to the sum of both PLDVs.

There were no changes to monitor output as a result of year 2000 considerations because monitor uses 4-byte date fields.

Effects on Accounting Data

None of the VM/ESA 2.2.0 performance changes are expected to have a significant effect on the values reported in the virtual machine resource usage accounting record. There have been no changes to the format of this record.

VM Performance Products

VM Performance Reporting Facility 1.2.1 (VMPRF) will run on VM/ESA 2.2.0 with the same support as VM/ESA 2.1.0. APAR VM60882 (PTF UM28172) is recommended. VMPRF at this service level includes the following enhancements:

- Support for >30000 userids.
- Four new reports.
  - MINIDISK_CACHE_USAGE_BY_TIME (PRF103) supplements the information provided by MINIDISK_CACHE_BY_TIME, including

information on MDC storage targets and usage for both main and expanded storage.

- SYSTEM_FACILITIES_BY_TIME (PRF104) contains CCW translation fast path performance information.
- SFS_BF_REQUESTS_BY_TIME (PRF105) contains a breakdown of SFS byte file system requests.
- USER_CONFIGURATION (PRF106) contains user virtual machine configuration and tuning settings by userid.

Realtime Monitor VM/ESA 1.5.2 (RTM/ESA) requires APAR GC05398 (PTF UG033816 to run on VM/ESA 2.2.0.

FCON/ESA Versions 2.3.00 and 2.3.01 will work with VM/ESA 2.1.0, but do not exploit any of the additional monitor data fields. With FCON/ESA Version 2.3.02, reports have been adapted to show the following new fields:

- Master-only PLDV statistics are included in PROCLOG display.
- Dispatch list priority and number of minor time slices used have been included in the ′Add to disp. list′ and ′Drop from disp. list′ entries respectively of the UTRANDET display.

Performance Analysis Facility/VM 1.1.3 (VMPAF) will run on VM/ESA 2.2.0 with the same support as VM/ESA 2.1.0.

Further information on these VM/ESA performance products can be found on the VM Performance page (http://www.vm.ibm.com/perf).

Measurement Information

This chapter discusses the types of processors used for measurements in the report, the levels 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
- 9672-R53

  To run as a 9672-R53, 5 processors were varied offline from the 9672-RX3 hardware configuration screen.
- 9121-480
- 9121-320

  To run as a 9121-320, one processor was varied offline from the 9121-480 hardware configuration screen.
- PC Server 520
- RS/6000 591

Software

A pre-GA (General Availability) level of VM/ESA 2.2.0 was used for the measurements in this report.

Other VM/ESA releases were measured for this report. VM/ESA 2.1.0 was at the GA+first-RSU (Recommended Service Upgrade) level. The CP and CMS service levels were 9501 and 501 respectively. The service that was part of VM/ESA 2.1.0 after the first RSU level and integrated into VM/ESA 2.2.0 can account for some of the performance differences between VM/ESA 2.1.0 and VM/ESA 2.2.0.

See the appropriate workload section in Appendix A, “Workloads” on page 70 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 70.
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 processor engines used.

- Storage: The amount of real and expanded storage used on the processor.
  - Real: The amount of real storage used on the processor.
  - 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.

     IOB This type of think time distribution represents the think time defined by the IBM Office Benchmark (IOB V2.1) workload. The think time includes an average 2-second delay between commands issued by TPNS, the built-in think times that are part of the IOB scripts, and the IOB script scheduling algorithm. The average message rate per user stays constant across all of the measurements. See “IBM Office Benchmark (IOB)” on page 80 for more details.
   - CMS block size: The block size of the CMS minidisks.
   - Virtual Machines: The virtual machines used in the measurement.

     For each virtual machine, the table indicates the following: name, number used, type, size and mode, 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 10.

   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.

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

Tools Description

The primary tools used to collect and evaluate the performance measurements are listed below.

Licensed Programs:

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

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.

Internal Tools:

FSTTAPE Reduces hardware monitor data for the 9121 processors.

Hardware Monitor Collects processor event and timing data.

REDFP Consolidates the QUERY FILEPOOL STATUS data from SFS measurements.

Migration from VM/ESA 2.1.0

This chapter explores the performance effects of migrating from VM/ESA 2.1.0 to VM/ESA 2.2.0. The following environments were measured: CMS-intensive, OfficeVision, and VSE guest.

CMS-Intensive

VM/ESA 2.2.0 shows improved internal throughput rates (ITR) and response times for the four measured CMS-intensive environments. ITR improvements ranging from 0.8% to 2.2% and response time improvements from 2% to 17% were observed. These improvements are primarily due to:

- Reduced CMS Working Set Size

  CMS working set size decreased by 3 to 7 pages for the measured environments, while total pages per user decreased by 24 to 30 pages as a result of this improvement.
- EXEC 2 to Compiled REXX Conversion
- GCS Preallocated Storage

For more information on these and other performance-related enhancements in VM/ESA 2.2.0, see “Performance Improvements” on page 3.

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

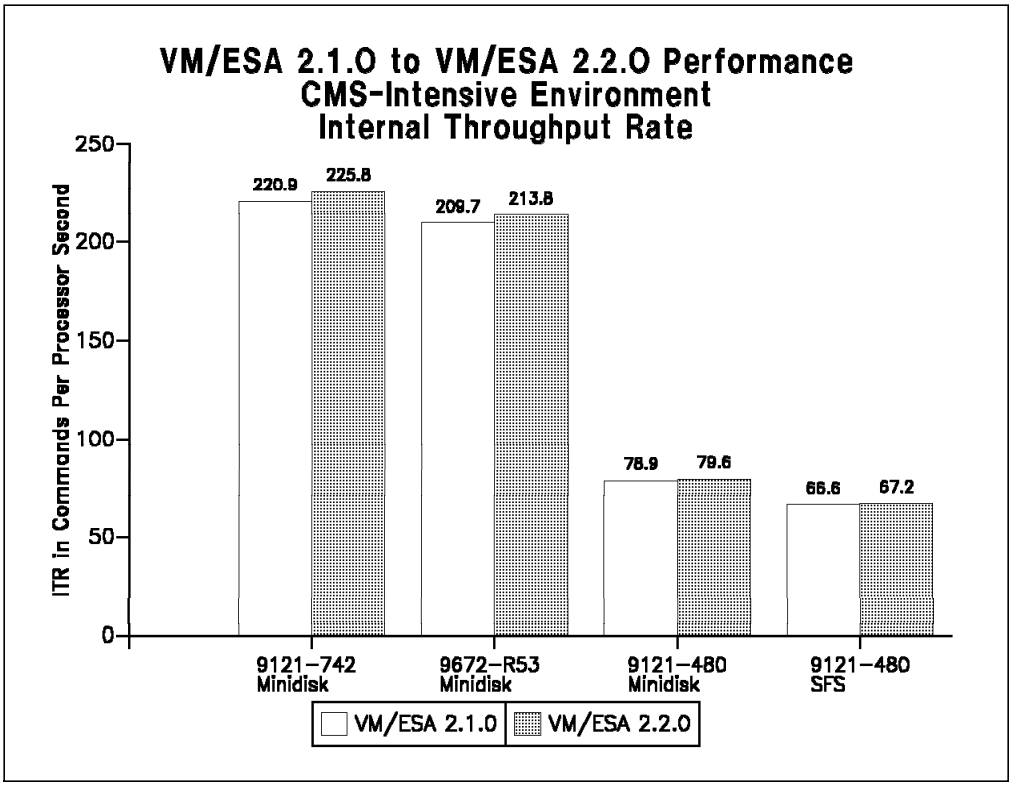

Figure 1. Internal throughput rate for the measured CMS-intensive environments

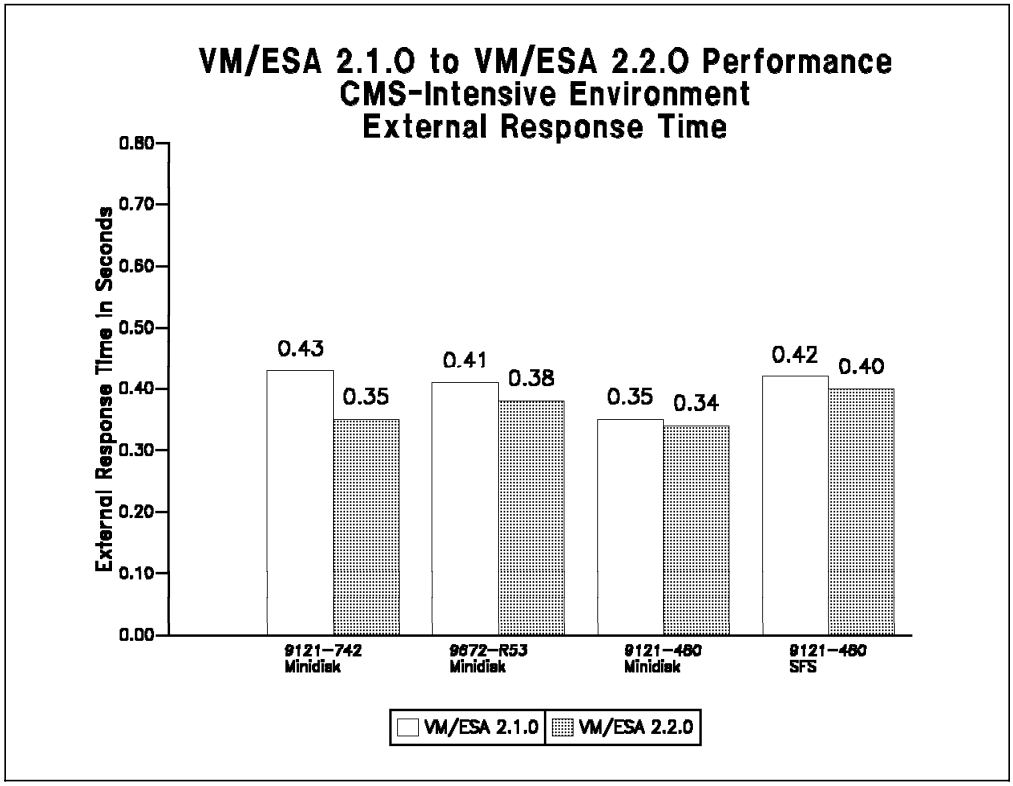

Figure 2. External response time for the measured CMS-intensive environments

Measurement results and discussion for each of these four environments are provided in the following sections.

9121-742 / Minidisk

Workload: FS8F0R

Hardware Configuration

Processor model: 9121-742
Processors used: 4
Storage:
  Real: 1024MB (default MDC)
  Expanded: 1024MB (MDC BIAS 0.1[2])
Tape: 3480 (Monitor)

DASD:

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

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

Communications:

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

Software Configuration

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

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	32MB/XA	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	5700	Users	3MB/XC	100		

Measurement Discussion: The following table shows that VM/ESA 2.2.0 has improved performance characteristics relative to VM/ESA 2.1.0. The key indicators of external response time (AVG LAST(T)) and internal throughput rate (ITR(H)) both improved. External response time improved by 17%, while internal throughput improved by 2.2%.

The CMS working set size improvement has resulted in significant changes in this environment's paging characteristics. The overall trend is a decrease in DASD paging, partially offset by a smaller increase in expanded storage paging.

[2] These measurements were made with an MDC BIAS value of 0.1 for expanded storage (using the SET MDCACHE command). Previous measurements have shown that this improves overall system performance for this particular environment. For more information refer to the *VM/ESA Release 2.2 Performance Report*.

DASD paging (PAGE/CMD) decreased by 28%, while expanded storage paging (XSTOR/CMD) increased by 4%. The CMS working set size improvement also resulted in a 30 pages per user decrease in DASD page slot requirements (as indicated by TOT PAGES/USER (V)). This decrease occurs during the execution of SYSPROF EXEC during ipl CMS and should improve the performance of scenarios where large numbers of CMS or OfficeVision users are logging on during a short period of time.

Of the four measured CMS-intensive environments, the two having expanded storage (this 9121-742 environment and the following 9672-R53 environment) showed the largest improvements. The removal of large numbers of one-time referenced pages has freed up much space in expanded storage that can now be used for other pages that are more likely to be re-referenced, as evidenced by the 45% increase in the expanded storage pagein rate (XSTOR IN/SEC). This improved expanded storage efficiency allows for a higher percentage of total paging to be done in expanded storage, resulting in a sizable DASD page I/O reduction. Since this advantage does not apply to the measured environments that have no expanded storage, they show smaller improvements relative to VM/ESA 2.1.0.

The true transaction rate (ETR(T)) measured by TPNS is essentially unchanged between the two measurements, while the transaction rate measured by CP (ETR) has decreased by 5%. This decrease is probably a side effect of the working set size improvement. The increased expanded storage efficiency means that CP can more quickly bring a paged-out dormant user back into real storage and onto the dispatch list. This has the effect of causing CP to count some cases that were counted as two separate transactions to now be counted as one transaction. This effect only shows up in the two measured CMS environments that have expanded storage.

CMS now uses diagnose X270 (pseudo timer extended) instead of diagnose X0C (pseudo timer) in many cases. Diagnose X270 provides the same information as diagnose X0C plus two additional date formats that represent the year as four digits. This change is reflected in the DIAG 0C/CMD and DIAG 270/CMD values in the results table.

Table 1 (Page 1 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.1.0 on the 9121-742

Release / **Run ID**	**2.1.0** **S48E5700**	**2.2.0** **S49E5702**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5700**	**5700**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Response Time				
TRIV INT	0.104	0.105	0.001	0.96%
NONTRIV INT	0.377	0.350	-0.027	-7.16%
TOT INT	0.258	0.250	-0.008	-3.10%
TOT INT ADJ	0.268	0.247	-0.021	-7.96%
AVG FIRST (T)	0.327	0.266	-0.061	-18.67%
AVG LAST (T)	0.425	0.353	-0.072	-16.93%
Throughput				
AVG THINK (T)	26.13	26.13	0.00	-0.01%
ETR	207.20	196.89	-10.31	-4.98%
ETR (T)	199.39	199.46	0.08	0.04%
ETR RATIO	1.039	0.987	-0.052	-5.01%
ITR (H)	220.92	225.84	4.92	2.22%
ITR	57.51	55.81	-1.70	-2.96%
EMUL ITR	91.20	87.49	-3.71	-4.07%
ITRR (H)	1.000	1.022	0.022	2.22%
ITRR	1.000	0.970	-0.030	-2.96%
Proc. Usage				
PBT/CMD (H)	18.106	17.712	-0.394	-2.18%
PBT/CMD	18.106	17.698	-0.408	-2.25%
CP/CMD (H)	7.106	6.842	-0.264	-3.72%
CP/CMD	6.721	6.417	-0.303	-4.51%
EMUL/CMD (H)	11.000	10.870	-0.130	-1.18%
EMUL/CMD	11.385	11.280	-0.105	-0.92%
Processor Util.				
TOTAL (H)	361.01	353.29	-7.72	-2.14%
TOTAL	361.00	353.00	-8.00	-2.22%
UTIL/PROC (H)	90.25	88.32	-1.93	-2.14%
UTIL/PROC	90.25	88.25	-2.00	-2.22%
TOTAL EMUL (H)	219.32	216.81	-2.50	-1.14%
TOTAL EMUL	227.00	225.00	-2.00	-0.88%
MASTER TOTAL (H)	92.06	90.15	-1.91	-2.07%
MASTER TOTAL	92.00	90.00	-2.00	-2.17%
MASTER EMUL (H)	35.03	35.36	0.34	0.96%
MASTER EMUL	37.00	37.00	0.00	0.00%
TVR(H)	1.65	1.63	-0.02	-1.01%
TVR	1.59	1.57	-0.02	-1.35%
Storage				
NUCLEUS SIZE (V)	2764KB	2808KB	44KB	1.59%
TRACE TABLE (V)	650KB	650KB	0KB	0.00%
WKSET (V)	82	75	-7	-8.54%
PGBLPGS	231K	231K	0K	0.00%
PGBLPGS/USER	40.5	40.5	0.0	0.00%
TOT PAGES/USER (V)	219	189	-30	-13.70%
FREEPGS	16568	16953	385	2.32%
FREE UTIL	0.93	0.92	-0.01	-0.64%
SHRPGS	1949	1973	24	1.23%

Table 1 (Page 2 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.1.0 on the 9121-742

Release **Run ID**	**2.1.0** **S48E5700**	**2.2.0** **S49E5702**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5700**	**5700**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Paging				
READS/SEC	1306	952	-354	-27.11%
WRITES/SEC	946	669	-277	-29.28%
PAGE/CMD	11.295	8.127	-3.168	-28.05%
PAGE IO RATE (V)	383.900	254.700	-129.200	-33.65%
PAGE IO/CMD (V)	1.925	1.277	-0.648	-33.68%
XSTOR IN/SEC	470	681	211	44.89%
XSTOR OUT/SEC	1642	1516	-126	-7.67%
XSTOR/CMD	10.593	11.015	0.422	3.98%
FAST CLR/CMD	8.446	8.794	0.348	4.12%
Queues				
DISPATCH LIST	125.33	108.55	-16.77	-13.38%
ELIGIBLE LIST	0.11	0.00	-0.11	-100.00%
I/O				
VIO RATE	1849	1815	-34	-1.84%
VIO/CMD	9.273	9.099	-0.174	-1.88%
RIO RATE (V)	767	639	-128	-16.69%
RIO/CMD (V)	3.847	3.204	-0.643	-16.72%
NONPAGE RIO/CMD (V)	1.921	1.927	0.005	0.27%
DASD RESP TIME (V)	20.800	20.400	-0.400	-1.92%
MDC REAL SIZE (MB)	31.2	32.0	0.8	2.56%
MDC XSTOR SIZE (MB)	63.9	63.0	-0.8	-1.27%
MDC READS (I/Os)	608	570	-38	-6.25%
MDC WRITES (I/Os)	27	27	0	0.00%
MDC AVOID	567	531	-36	-6.35%
MDC HIT RATIO	0.93	0.93	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	20.377	20.479	0.102	0.50%
DIAG/CMD	23.669	24.430	0.761	3.22%
DIAG 04/CMD	0.889	0.904	0.014	1.61%
DIAG 08/CMD	0.748	0.736	-0.013	-1.69%
DIAG 0C/CMD	1.134	0.212	-0.922	-81.33%
DIAG 14/CMD	0.025	0.025	0.000	-0.39%
DIAG 58/CMD	1.249	1.249	0.000	-0.01%
DIAG 98/CMD	0.278	0.298	0.020	7.10%
DIAG A4/CMD	3.734	3.559	-0.174	-4.67%
DIAG A8/CMD	2.670	2.664	-0.006	-0.21%
DIAG 214/CMD	11.675	12.485	0.810	6.94%
DIAG 270/CMD	0.000	0.917	0.917	na
SIE/CMD	55.170	55.148	-0.021	-0.04%
SIE INTCPT/CMD	37.515	36.949	-0.566	-1.51%
FREE TOTL/CMD	44.848	44.891	0.043	0.10%
VTAM Machines				
WKSET (V)	4140	4140	0	0.00%
TOT CPU/CMD (V)	2.9563	2.8688	-0.0875	-2.96%
CP CPU/CMD (V)	1.2817	1.2617	-0.0200	-1.56%
VIRT CPU/CMD (V)	1.6746	1.6071	-0.0675	-4.03%
DIAG 98/CMD (V)	0.278	0.297	0.018	6.55%

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

9672-R53 / Minidisk

Workload: FS8F0R

Hardware Configuration

Processor model: 9672-R53
Processors used: 5
Storage:
  Real: 1024MB (default MDC)
  Expanded: 1024MB (MDC BIAS 0.1)[2]
Tape: 3480 (Monitor)

DASD:

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

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

Communications:

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

Software Configuration

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

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	32MB/XA	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	5400	Users	3MB/XC	100		

Measurement Discussion: These results are very similar to the 9121-742 results discussed in the previous section. External response time (AVG LAST(T)) improved by 8%, while internal throughput (ITR(V)) improved by 2.0%. The hardware configuration used for these 9672-R53 measurements is identical to the configuration used for the 9121-742 measurements except for the difference in processor model. As a result, the same performance improvements that affected the 9121-742 comparison also affected this comparison, and to about the same extent.

Hardware instrumentation was not available on the measured 9672 processor so the ITR calculation is based on the average processor utilization reported by VMPRF. Results on processors that do have instrumentation consistently show a very close correspondence between total processor utilization reported by

hardware instrumentation, CP monitor (as reported by VMPRF), and RTM VM/ESA.

Table 2 (Page 1 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.1.0 on the 9672-R53

Release **Run ID**	**2.1.0** **C58E5400**	**2.2.0** **C59E5400**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5400**	**5400**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**5**	**5**		
Response Time				
TRIV INT	0.098	0.095	-0.003	-3.06%
NONTRIV INT	0.360	0.342	-0.018	-5.00%
TOT INT	0.245	0.236	-0.009	-3.67%
TOT INT ADJ	0.257	0.242	-0.015	-5.77%
AVG FIRST (T)	0.314	0.286	-0.027	-8.71%
AVG LAST (T)	0.409	0.378	-0.031	-7.65%
Throughput				
AVG THINK (T)	26.15	26.14	-0.01	-0.03%
ETR	198.64	194.21	-4.43	-2.23%
ETR (T)	189.32	189.21	-0.11	-0.06%
ETR RATIO	1.049	1.026	-0.023	-2.17%
ITR (V)	209.66	213.80	4.14	1.97%
ITR	43.96	43.85	-0.11	-0.25%
EMUL ITR	68.72	67.58	-1.15	-1.67%
ITRR (V)	1.000	1.020	0.020	1.97%
ITRR	1.000	0.997	-0.003	-0.25%
Proc. Usage				
PBT/CMD	23.875	23.466	-0.409	-1.71%
CP/CMD	8.610	8.245	-0.365	-4.24%
EMUL/CMD	15.265	15.221	-0.044	-0.29%
Processor Util.				
TOTAL	452.00	444.00	-8.00	-1.77%
UTIL/PROC	90.40	88.80	-1.60	-1.77%
UTIL/PROC (V)	90.30	88.50	-1.80	-1.99%
TOTAL EMUL	289.00	288.00	-1.00	-0.35%
TVR	1.56	1.54	-0.02	-1.43%
Storage				
NUCLEUS SIZE (V)	2764KB	2808KB	44KB	1.59%
TRACE TABLE (V)	800KB	800KB	0KB	0.00%
WKSET (V)	77	71	-6	-7.79%
PGBLPGS	231K	231K	0K	0.00%
PGBLPGS/USER	42.8	42.8	0.0	0.00%
TOT PAGES/USER (V)	214	185	-29	-13.55%
FREEPGS	15674	16060	386	2.46%
FREE UTIL	0.91	0.92	0.01	1.08%
SHRPGS	1922	1885	-37	-1.93%
Paging				
READS/SEC	887	639	-248	-27.96%
WRITES/SEC	591	403	-188	-31.81%
PAGE/CMD	7.807	5.507	-2.300	-29.46%
PAGE IO RATE (V)	238.800	152.600	-86.200	-36.10%
PAGE IO/CMD (V)	1.261	0.807	-0.455	-36.06%
XSTOR IN/SEC	717	834	117	16.32%
XSTOR OUT/SEC	1481	1362	-119	-8.04%
XSTOR/CMD	11.610	11.606	-0.004	-0.03%
FAST CLR/CMD	8.430	8.773	0.343	4.07%
Queues				
DISPATCH LIST	105.87	99.30	-6.57	-6.21%
ELIGIBLE LIST	0.00	0.02	0.02	na

Table 2 (Page 2 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.1.0 on the 9672-R53

Release / **Run ID**	**2.1.0** **C58E5400**	**2.2.0** **C59E5400**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5400**	**5400**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**5**	**5**		
I/O				
VIO RATE	1757	1721	-36	-2.05%
VIO/CMD	9.280	9.096	-0.185	-1.99%
RIO RATE (V)	603	518	-85	-14.10%
RIO/CMD (V)	3.185	2.738	-0.447	-14.04%
NONPAGE RIO/CMD (V)	1.924	1.931	0.007	0.39%
DASD RESP TIME (V)	20.000	20.000	0.000	0.00%
MDC REAL SIZE (MB)	31.6	34.3	2.7	8.60%
MDC XSTOR SIZE (MB)	63.6	63.2	-0.4	-0.68%
MDC READS (I/Os)	579	540	-39	-6.74%
MDC WRITES (I/Os)	26	26	0	0.00%
MDC AVOID	539	503	-36	-6.68%
MDC HIT RATIO	0.93	0.93	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	20.036	20.135	0.099	0.49%
DIAG/CMD	23.681	24.407	0.726	3.07%
DIAG 04/CMD	0.879	0.896	0.017	1.96%
DIAG 08/CMD	0.748	0.736	-0.011	-1.50%
DIAG 0C/CMD	1.133	0.212	-0.921	-81.32%
DIAG 14/CMD	0.025	0.025	0.000	-0.12%
DIAG 58/CMD	1.250	1.249	0.000	-0.04%
DIAG 98/CMD	0.275	0.294	0.019	6.93%
DIAG A4/CMD	3.756	3.553	-0.203	-5.41%
DIAG A8/CMD	2.670	2.668	-0.002	-0.06%
DIAG 214/CMD	11.679	12.479	0.800	6.85%
DIAG 270/CMD	na	0.921		
SIE/CMD	51.885	52.545	0.660	1.27%
SIE INTCPT/CMD	36.838	36.781	-0.057	-0.15%
FREE TOTL/CMD	44.728	44.897	0.169	0.38%
VTAM Machines				
WKSET (V)	4125	4118	-7	-0.17%
TOT CPU/CMD (V)	4.3464	4.2369	-0.1095	-2.52%
CP CPU/CMD (V)	1.9621	1.9115	-0.0506	-2.58%
VIRT CPU/CMD (V)	2.3844	2.3255	-0.0589	-2.47%
DIAG 98/CMD (V)	0.274	0.294	0.019	7.00%

Note: T=TPNS, V=VMPRF, 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*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2	16 R		

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

Communications:

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

Software Configuration

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

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	32MB/XA	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	560	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	2000	Users	3MB/XC	100		

Measurement Discussion: External response time (AVG LAST(T) improved by 2%, while internal throughput (ITR(H)) improved by 0.9%. The degree of improvement is less than was observed for the 9121-742 and 9672-R53 environments discussed in the last two sections. This is due to the fact that the CMS working set size improvement has a smaller effect in this (no expanded storage) environment. Working set (WKSET (V)) decreased by 3 pages, while total pages per user (TOT PAGES/USER (V)) decreased by 24 pages.

See “CMS-Intensive” on page 11 and “9121-742 / Minidisk” on page 13 for a discussion of results that are common to all four measured CMS-intensive environments.

Table 3 (Page 1 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.1.0 on the 9121-480

Release / **Run ID**	**2.1.0** **L28E2001**	**2.2.0** **L29E2005**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2000**	**2000**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.130	0.130	0.000	0.00%
NONTRIV INT	0.415	0.408	-0.007	-1.69%
TOT INT	0.320	0.315	-0.005	-1.56%
TOT INT ADJ	0.283	0.278	-0.005	-1.63%
AVG FIRST (T)	0.244	0.242	-0.003	-1.23%
AVG LAST (T)	0.348	0.340	-0.008	-2.16%
Throughput				
AVG THINK (T)	26.19	26.19	0.01	0.04%
ETR	62.01	62.05	0.04	0.06%
ETR (T)	70.18	70.27	0.10	0.14%
ETR RATIO	0.884	0.883	-0.001	-0.07%
ITR (H)	78.90	79.62	0.72	0.91%
ITR	34.89	35.18	0.29	0.83%
EMUL ITR	51.85	52.36	0.52	1.00%
ITRR (H)	1.000	1.009	0.009	0.91%
ITRR	1.000	1.008	0.008	0.83%
Proc. Usage				
PBT/CMD (H)	25.350	25.120	-0.230	-0.91%
PBT/CMD	25.365	25.046	-0.320	-1.26%
CP/CMD (H)	8.870	8.835	-0.034	-0.39%
CP/CMD	8.265	8.111	-0.154	-1.86%
EMUL/CMD (H)	16.480	16.285	-0.195	-1.19%
EMUL/CMD	17.100	16.934	-0.166	-0.97%
Processor Util.				
TOTAL (H)	177.89	176.52	-1.37	-0.77%
TOTAL	178.00	176.00	-2.00	-1.12%
UTIL/PROC (H)	88.95	88.26	-0.69	-0.77%
UTIL/PROC	89.00	88.00	-1.00	-1.12%
TOTAL EMUL (H)	115.65	114.43	-1.21	-1.05%
TOTAL EMUL	120.00	119.00	-1.00	-0.83%
MASTER TOTAL (H)	88.53	87.75	-0.78	-0.88%
MASTER TOTAL	89.00	88.00	-1.00	-1.12%
MASTER EMUL (H)	50.86	50.08	-0.78	-1.53%
MASTER EMUL	53.00	52.00	-1.00	-1.89%
TVR(H)	1.54	1.54	0.00	0.28%
TVR	1.48	1.48	0.00	-0.29%
Storage				
NUCLEUS SIZE (V)	2764KB	2804KB	40KB	1.45%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	86	83	-3	-3.49%
PGBLPGS	54455	54290	-165	-0.30%
PGBLPGS/USER	27.2	27.1	-0.1	-0.30%
TOT PAGES/USER (V)	196	172	-24	-12.24%
FREEPGS	5810	5986	176	3.03%
FREE UTIL	0.93	0.94	0.02	1.68%
SHRPGS	1409	1510	101	7.17%

Table 3 (Page 2 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.1.0 on the 9121-480

Release **Run ID**	**2.1.0** **L28E2001**	**2.2.0** **L29E2005**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2000**	**2000**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	724	700	-24	-3.31%
WRITES/SEC	474	458	-16	-3.38%
PAGE/CMD	17.072	16.479	-0.593	-3.47%
PAGE IO RATE (V)	198.800	190.500	-8.300	-4.18%
PAGE IO/CMD (V)	2.833	2.711	-0.122	-4.31%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	8.379	8.752	0.373	4.45%
Queues				
DISPATCH LIST	40.74	44.30	3.56	8.73%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	706	693	-13	-1.84%
VIO/CMD	10.061	9.862	-0.199	-1.98%
RIO RATE (V)	394	386	-8	-2.03%
RIO/CMD (V)	5.615	5.493	-0.122	-2.17%
NONPAGE RIO/CMD (V)	2.782	2.782	0.000	0.02%
DASD RESP TIME (V)	20.000	19.900	-0.100	-0.50%
MDC REAL SIZE (MB)	39.9	41.5	1.5	3.78%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	214	201	-13	-6.07%
MDC WRITES (I/Os)	9.66	9.79	0.13	1.35%
MDC AVOID	202	189	-13	-6.44%
MDC HIT RATIO	0.94	0.94	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.912	13.896	-0.016	-0.12%
DIAG/CMD	25.926	26.668	0.742	2.86%
DIAG 04/CMD	2.367	2.364	-0.003	-0.12%
DIAG 08/CMD	0.748	0.733	-0.015	-2.04%
DIAG 0C/CMD	1.134	0.212	-0.922	-81.34%
DIAG 14/CMD	0.024	0.024	0.000	-0.52%
DIAG 58/CMD	1.249	1.249	0.001	0.05%
DIAG 98/CMD	1.054	1.058	0.004	0.38%
DIAG A4/CMD	3.759	3.569	-0.190	-5.04%
DIAG A8/CMD	2.684	2.667	-0.017	-0.63%
DIAG 214/CMD	11.638	12.491	0.854	7.34%
DIAG 270/CMD	0.000	0.911	0.911	na
SIE/CMD	52.725	52.752	0.027	0.05%
SIE INTCPT/CMD	34.799	34.289	-0.510	-1.46%
FREE TOTL/CMD	49.419	49.508	0.089	0.18%
VTAM Machines				
WKSET (V)	541	552	11	2.03%
TOT CPU/CMD (V)	4.0454	3.9055	-0.1399	-3.46%
CP CPU/CMD (V)	1.4725	1.4468	-0.0257	-1.75%
VIRT CPU/CMD (V)	2.5729	2.4587	-0.1142	-4.44%
DIAG 98/CMD (V)	1.054	1.058	0.004	0.41%

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*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	16	6	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4		2	2		16 R	

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

Communications:

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

Software Configuration

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

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
CRRSERV1	1	SFS	16MB/XC	100		
ROSERV1	1	SFS	64MB/XC	100		QUICKDSP ON
RWSERVn	2	SFS	64MB/XC	1500	1300	QUICKDSP ON
SMART	1	RTM	32MB/XA	3%	400	QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	512	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1720	Users	3MB/XC	100		

Measurement Discussion: External response time (AVG LAST(T) improved by 5%, while internal throughput (ITR(H)) improved by 0.8%. These improvements are similar to those observed for the 9121-480 minidisk environment (see the previous section). SFS server performance is essentially unchanged from VM/ESA 2.1.0.

See “CMS-Intensive” on page 11 and “9121-742 / Minidisk” on page 13 for a discussion of results that are common to all four measured CMS-intensive environments.

Table 4 (Page 1 of 3). SFS CMS-intensive migration from VM/ESA 2.1.0 on the 9121-480

Release	**2.1.0**	**2.2.0**		
Run ID	**L28S1720**	**L29S1721**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.127	0.125	-0.002	-1.57%
NONTRIV INT	0.513	0.485	-0.028	-5.46%
TOT INT	0.384	0.364	-0.020	-5.21%
TOT INT ADJ	0.344	0.324	-0.020	-5.82%
AVG FIRST (T)	0.272	0.263	-0.008	-3.13%
AVG LAST (T)	0.421	0.399	-0.021	-5.11%
Throughput				
AVG THINK (T)	26.25	26.17	-0.08	-0.30%
ETR	53.91	53.66	-0.25	-0.46%
ETR (T)	60.21	60.32	0.11	0.18%
ETR RATIO	0.895	0.890	-0.006	-0.64%
ITR (H)	66.64	67.19	0.54	0.82%
ITR	29.85	29.91	0.06	0.19%
EMUL ITR	45.16	45.28	0.12	0.26%
ITRR (H)	1.000	1.008	0.008	0.82%
ITRR	1.000	1.002	0.002	0.19%
Proc. Usage				
PBT/CMD (H)	30.010	29.767	-0.243	-0.81%
PBT/CMD	30.062	29.677	-0.385	-1.28%
CP/CMD (H)	10.823	10.769	-0.053	-0.49%
CP/CMD	10.297	9.947	-0.350	-3.40%
EMUL/CMD (H)	19.187	18.998	-0.189	-0.99%
EMUL/CMD	19.764	19.729	-0.035	-0.18%
Processor Util.				
TOTAL (H)	180.69	179.55	-1.14	-0.63%
TOTAL	181.00	179.00	-2.00	-1.10%
UTIL/PROC (H)	90.34	89.77	-0.57	-0.63%
UTIL/PROC	90.50	89.50	-1.00	-1.10%
TOTAL EMUL (H)	115.52	114.59	-0.94	-0.81%
TOTAL EMUL	119.00	119.00	0.00	0.00%
MASTER TOTAL (H)	90.04	89.58	-0.46	-0.51%
MASTER TOTAL	90.00	90.00	0.00	0.00%
MASTER EMUL (H)	51.95	51.62	-0.34	-0.64%
MASTER EMUL	54.00	54.00	0.00	0.00%
TVR(H)	1.56	1.57	0.00	0.18%
TVR	1.52	1.50	-0.02	-1.10%
Storage				
NUCLEUS SIZE (V)	2764KB	2804KB	40KB	1.45%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	82	79	-3	-3.66%
PGBLPGS	55499	55356	-143	-0.26%
PGBLPGS/USER	32.3	32.2	-0.1	-0.26%
TOT PAGES/USER (V)	181	157	-24	-13.26%
FREEPGS	5118	5242	124	2.42%
FREE UTIL	0.95	0.93	-0.02	-2.37%
SHRPGS	1656	1696	40	2.42%

Table 4 (Page 2 of 3). SFS CMS-intensive migration from VM/ESA 2.1.0 on the 9121-480

Release **Run ID**	**2.1.0** **L28S1720**	**2.2.0** **L29S1721**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	592	579	-13	-2.20%
WRITES/SEC	401	389	-12	-2.99%
PAGE/CMD	16.492	16.049	-0.444	-2.69%
PAGE IO RATE (V)	155.600	150.700	-4.900	-3.15%
PAGE IO/CMD (V)	2.584	2.498	-0.086	-3.32%
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	8.205	8.538	0.334	4.07%
Queues				
DISPATCH LIST	43.27	45.39	2.12	4.90%
ELIGIBLE LIST	0.02	0.00	-0.02	-100.00%
I/O				
VIO RATE	607	598	-9	-1.48%
VIO/CMD	10.081	9.914	-0.167	-1.66%
RIO RATE (V)	354	352	-2	-0.56%
RIO/CMD (V)	5.879	5.836	-0.044	-0.74%
NONPAGE RIO/CMD (V)	3.295	3.337	0.042	1.28%
DASD RESP TIME (V)	18.500	18.400	-0.100	-0.54%
MDC REAL SIZE (MB)	65.8	67.3	1.4	2.16%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	170	159	-11	-6.47%
MDC WRITES (I/Os)	15	15	0	0.00%
MDC AVOID	144	133	-11	-7.64%
MDC HIT RATIO	0.84	0.83	-0.01	-1.19%
PRIVOPs				
PRIVOP/CMD	20.604	20.471	-0.133	-0.65%
DIAG/CMD	23.918	24.836	0.918	3.84%
DIAG 04/CMD	2.532	2.574	0.042	1.64%
DIAG 08/CMD	0.747	0.736	-0.010	-1.40%
DIAG 0C/CMD	1.155	0.238	-0.917	-79.37%
DIAG 14/CMD	0.024	0.025	0.000	0.50%
DIAG 58/CMD	1.249	1.250	0.001	0.07%
DIAG 98/CMD	1.184	1.193	0.008	0.72%
DIAG A4/CMD	2.131	1.975	-0.156	-7.31%
DIAG A8/CMD	2.469	2.487	0.018	0.74%
DIAG 214/CMD	11.165	12.066	0.901	8.07%
DIAG 270/CMD	0.000	0.912	0.912	na
SIE/CMD	59.924	60.050	0.126	0.21%
SIE INTCPT/CMD	41.348	41.434	0.087	0.21%
FREE TOTL/CMD	52.716	52.738	0.022	0.04%
VTAM Machines				
WKSET (V)	505	506	1	0.20%
TOT CPU/CMD (V)	4.2075	4.0711	-0.1364	-3.24%
CP CPU/CMD (V)	1.5317	1.5105	-0.0212	-1.38%
VIRT CPU/CMD (V)	2.6758	2.5606	-0.1152	-4.31%
DIAG 98/CMD (V)	1.184	1.193	0.008	0.72%

Table 4 (Page 3 of 3). SFS CMS-intensive migration from VM/ESA 2.1.0 on the 9121-480

Release Run ID	2.1.0 L28S1720	2.2.0 L29S1721	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
SFS Servers				
WKSET (V)	3376	3394	18	0.53%
TOT CPU/CMD (V)	3.4140	3.4079	-0.0061	-0.18%
CP CPU/CMD (V)	1.4856	1.4921	0.0065	0.44%
VIRT CPU/CMD (V)	1.9285	1.9158	-0.0127	-0.66%
FP REQ/CMD(Q)	1.149	1.120	-0.029	-2.52%
IO/CMD (Q)	1.628	1.589	-0.039	-2.40%
IO TIME/CMD (Q)	0.026	0.025	-0.001	-3.85%
SFS TIME/CMD (Q)	0.044	0.042	-0.002	-4.55%

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 *VM/ESA: CMS File Pool 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. SFS CMS-intensive migration from VM/ESA 2.1.0 on the 9121-480

Release / **Run ID**	**2.1.0** **L28S1720**	**2.2.0** **L29S1721**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Close File Requests	0.3646	0.3640	-0.0006	-0.16%
Commit Requests	0.0164	0.0164	0.0000	0.00%
Connect Requests	0.0079	0.0078	-0.0001	-1.27%
Delete File Requests	0.0739	0.0739	0.0000	0.00%
Lock Requests	0.0247	0.0248	0.0001	0.40%
Open File New Requests	0.0033	0.0033	0.0000	0.00%
Open File Read Requests	0.2188	0.2184	-0.0004	-0.18%
Open File Replace Requests	0.1213	0.1210	-0.0003	-0.25%
Open File Write Requests	0.0214	0.0213	-0.0001	-0.47%
Query File Pool Requests	0.0000	0.0000	0.0000	na
Query User Space Requests	0.0214	0.0213	-0.0001	-0.47%
Read File Requests	0.1723	0.1444	-0.0279	-16.19%
Refresh Directory Requests	0.0229	0.0227	-0.0002	-0.87%
Rename Requests	0.0049	0.0049	0.0000	0.00%
Unlock Requests	0.0246	0.0247	0.0001	0.41%
Write File Requests	0.0507	0.0505	-0.0002	-0.39%
Total File Pool Requests	1.1489	1.1195	-0.0294	-2.56%
File Pool Request Service Time	43.7161	41.8526	-1.8635	-4.26%
Local File Pool Requests	1.1489	1.1195	-0.0294	-2.56%
Begin LUWs	0.4450	0.4487	0.0037	0.83%
Agent Holding Time (msec)	134.8440	117.5462	-17.2978	-12.83%
SAC Calls	5.4844	5.5302	0.0458	0.84%
Catalog Lock Conflicts	0.0015	0.0020	0.0005	33.33%
Total Lock Conflicts	0.0015	0.0020	0.0005	33.33%
Lock Wait Time (msec)	0.0959	0.1731	0.0772	80.50%
File Blocks Read	0.9034	0.9031	-0.0003	-0.03%
File Blocks Written	0.4981	0.4970	-0.0011	-0.22%
Catalog Blocks Read	0.5132	0.5049	-0.0083	-1.62%
Catalog Blocks Written	0.2649	0.2626	-0.0023	-0.87%
Control Minidisk Blocks Written	0.0510	0.0509	-0.0001	-0.20%
Log Blocks Written	0.4505	0.4608	0.0103	2.29%
Total DASD Block Transfers	2.6811	2.6793	-0.0018	-0.07%
BIO Requests to Read File Block	0.4186	0.3896	-0.0290	-6.93%
BIO Requests to Write File Blocks	0.1797	0.1794	-0.0003	-0.17%
BIO Requests to Read Catalog Blks	0.5132	0.5049	-0.0083	-1.62%
BIO Requests to Write Catalog Blks	0.2161	0.2129	-0.0032	-1.48%
BIO Requests to Write Ctl Mdsk Blks	0.0021	0.0020	-0.0001	-4.76%
BIO Requests to Write Log Blocks	0.3919	0.4018	0.0099	2.53%
Total BIO Requests	1.7215	1.6906	-0.0309	-1.79%
Total BIO Request Time (msec)	25.6029	25.0090	-0.5939	-2.32%
I/O Requests to Read File Blocks	0.3007	0.2641	-0.0366	-12.17%
I/O Requests to Write File Blocks	0.1947	0.1946	-0.0001	-0.05%
I/O Requests to Read Catalog Blks	0.5132	0.5049	-0.0083	-1.62%
I/O Requests to Write Catalog Blks	0.2229	0.2196	-0.0033	-1.48%
I/O Requests to Write Ctl Mdsk Blks	0.0039	0.0039	0.0000	0.00%
I/O Requests to Write Log Blocks	0.3922	0.4021	0.0099	2.52%
Total I/O Requests	1.6276	1.5892	-0.0384	-2.36%
Get Logname Requests	0.0033	0.0032	-0.0001	-3.03%
Get LUWID Requests	0.0033	0.0032	-0.0001	-3.03%
Total CRR Requests	0.0065	0.0065	0.0000	0.00%
CRR Request Service Time (msec)	0.0810	0.0771	-0.0039	-4.81%
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 2.1.0 on the 9121-480

Release **Run ID**	**2.1.0** **L28S1720**	**2.2.0** **L29S1721**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Agents Held	8.1	7.1	-1.0	-12.67%
Agents In-call	2.6	2.5	-0.1	-4.09%
Avg LUW Time (msec)	303.0	262.0	-41.0	-13.55%
Avg File Pool Request Time (msec)	38.1	37.4	-0.7	-1.75%
Avg Lock Wait Time (msec)	63.9	86.6	22.6	35.38%
SAC Calls / FP Request	4.77	4.94	0.17	3.48%
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)	828	840	12	1.45%
Checkpoints Taken (delta)	34	34	0	0.00%
Checkpoint Duration (sec)	3.7	3.6	-0.1	-3.78%
Seconds Between Checkpoints	56.8	56.8	0.0	0.00%
Checkpoint Utilization	6.6	6.3	-0.3	-3.81%
BIO Request Time (msec)	14.87	14.79	-0.08	-0.53%
Blocking Factor (Blocks/BIO)	1.56	1.58	0.03	1.76%
Chaining Factor (Blocks/IO)	1.65	1.69	0.04	2.35%

Note: Query Filepool Counters — derived results

Table 7 compares the VM/ESA 2.2.0 SFS measurement to the corresponding VM/ESA 2.2.0 minidisk-only measurement from Table 3 on page 21.

Table 7 (Page 1 of 3). Minidisk to SFS comparison for VM/ESA 2.2.0 on the 9121-480

File System **Release** **Run ID**	**Minidisk** **2.2.0** **L29E2005**	**SFS** **2.2.0** **L29S1721**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**00MB**		
Users	**2000**	**1720**	**-280**	**-14.00%**
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.130	0.125	-0.005	-3.85%
NONTRIV INT	0.408	0.485	0.077	18.87%
TOT INT	0.315	0.364	0.049	15.56%
TOT INT ADJ	0.278	0.324	0.046	16.42%
AVG FIRST (T)	0.242	0.263	0.022	8.90%
AVG LAST (T)	0.340	0.399	0.059	17.35%
Throughput				
AVG THINK (T)	26.19	26.17	-0.02	-0.08%
ETR	62.05	53.66	-8.39	-13.52%
ETR (T)	70.27	60.32	-9.95	-14.17%
ETR RATIO	0.883	0.890	0.007	0.75%
ITR (H)	79.62	67.19	-12.43	-15.61%
ITR	35.18	29.91	-5.27	-14.97%
EMUL ITR	52.36	45.28	-7.09	-13.53%
ITRR (H)	1.000	0.844	-0.156	-15.61%
ITRR	1.000	0.850	-0.150	-14.97%
Proc. Usage				
PBT/CMD (H)	25.120	29.767	4.647	18.50%
PBT/CMD	25.046	29.677	4.631	18.49%
CP/CMD (H)	8.835	10.769	1.934	21.89%
CP/CMD	8.111	9.947	1.836	22.64%
EMUL/CMD (H)	16.285	18.998	2.713	16.66%
EMUL/CMD	16.934	19.729	2.795	16.50%
Processor Util.				
TOTAL (H)	176.52	179.55	3.02	1.71%
TOTAL	176.00	179.00	3.00	1.70%
UTIL/PROC (H)	88.26	89.77	1.51	1.71%
UTIL/PROC	88.00	89.50	1.50	1.70%
TOTAL EMUL (H)	114.43	114.59	0.15	0.13%
TOTAL EMUL	119.00	119.00	0.00	0.00%
MASTER TOTAL (H)	87.75	89.58	1.83	2.08%
MASTER TOTAL	88.00	90.00	2.00	2.27%
MASTER EMUL (H)	50.08	51.62	1.54	3.08%
MASTER EMUL	52.00	54.00	2.00	3.85%
TVR(H)	1.54	1.57	0.02	1.58%
TVR	1.48	1.50	0.03	1.70%
Storage				
NUCLEUS SIZE (V)	2804KB	2804KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	83	79	-4	-4.82%
PGBLPGS	54290	55356	1066	1.96%
PGBLPGS/USER	27.1	32.2	5.0	18.56%
TOT PAGES/USER (V)	172	157	-15	-8.72%
FREEPGS	5986	5242	-744	-12.43%
FREE UTIL	0.94	0.93	-0.01	-1.38%
SHRPGS	1510	1696	186	12.32%

Table 7 (Page 2 of 3). Minidisk to SFS comparison for VM/ESA 2.2.0 on the 9121-480

File System Release Run ID	Minidisk 2.2.0 L29E2005	SFS 2.2.0 L29S1721	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**00MB**		
Users	**2000**	**1720**	**-280**	**-14.00%**
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	700	579	-121	-17.29%
WRITES/SEC	458	389	-69	-15.07%
PAGE/CMD	16.479	16.049	-0.430	-2.61%
PAGE IO RATE (V)	190.500	150.700	-39.800	-20.89%
PAGE IO/CMD (V)	2.711	2.498	-0.212	-7.84%
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	8.752	8.538	-0.213	-2.44%
Queues				
DISPATCH LIST	44.30	45.39	1.10	2.48%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	693	598	-95	-13.71%
VIO/CMD	9.862	9.914	0.053	0.53%
RIO RATE (V)	386	352	-34	-8.81%
RIO/CMD (V)	5.493	5.836	0.343	6.24%
NONPAGE RIO/CMD (V)	2.782	3.337	0.555	19.96%
DASD RESP TIME (V)	19.900	18.400	-1.500	-7.54%
MDC REAL SIZE (MB)	41.5	67.3	25.8	62.23%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	201	159	-42	-20.90%
MDC WRITES (I/Os)	9.79	15	5.21	53.22%
MDC AVOID	189	133	-56	-29.63%
MDC HIT RATIO	0.94	0.83	-0.11	-11.70%
PRIVOPs				
PRIVOP/CMD	13.896	20.471	6.575	47.31%
DIAG/CMD	26.668	24.836	-1.832	-6.87%
DIAG 04/CMD	2.364	2.574	0.210	8.88%
DIAG 08/CMD	0.733	0.736	0.003	0.45%
DIAG 0C/CMD	0.212	0.238	0.027	12.61%
DIAG 14/CMD	0.024	0.025	0.000	0.67%
DIAG 58/CMD	1.249	1.250	0.001	0.06%
DIAG 98/CMD	1.058	1.193	0.135	12.72%
DIAG A4/CMD	3.569	1.975	-1.594	-44.66%
DIAG A8/CMD	2.667	2.487	-0.180	-6.76%
DIAG 214/CMD	12.491	12.066	-0.425	-3.40%
DIAG 270/CMD	0.922	0.922	0.001	0.07%
SIE/CMD	52.752	60.050	7.297	13.83%
SIE INTCPT/CMD	34.289	41.434	7.145	20.84%
FREE TOTL/CMD	49.508	52.738	3.230	6.53%
VTAM Machines				
WKSET (V)	552	506	-46	-8.33%
TOT CPU/CMD (V)	3.9055	4.0711	0.1656	4.24%
CP CPU/CMD (V)	1.4468	1.5105	0.0637	4.40%
VIRT CPU/CMD (V)	2.4587	2.5606	0.1019	4.14%
DIAG 98/CMD (V)	1.058	1.193	0.135	12.71%

Table 7 (Page 3 of 3). Minidisk to SFS comparison for VM/ESA 2.2.0 on the 9121-480

File System Release Run ID	Minidisk 2.2.0 L29E2005	SFS 2.2.0 L29S1721	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**00MB**		
Users	**2000**	**1720**	**-280**	**-14.00%**
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
SFS Servers				
WKSET (V)	na	3394		
TOT CPU/CMD (V)	na	3.4079		
CP CPU/CMD (V)	na	1.4921		
VIRT CPU/CMD (V)	na	1.9158		
FP REQ/CMD(Q)	na	1.120		
IO/CMD (Q)	na	1.589		
IO TIME/CMD (Q)	na	0.025		
SFS TIME/CMD (Q)	na	0.042		

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

These results show the normal relationship between the minidisk and SFS implementations of the FS8F workload that has been observed in the past.

OfficeVision

This section examines the performance impact of migrating an OfficeVision 1.3.0 system from VM/ESA 2.1.0 to VM/ESA 2.2.0.

Workload: IOB

Hardware Configuration

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

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	6	4	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4	10	4	2	16 R	7 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: IOB
CMS block size: 4KB

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
Ennnn	10	Workload	2MB/XA	1000		QUICKDSP ON
PRNTEAT1	1	Workload	2MB/XA	1000		QUICKDSP ON
PROCAL	1	OV/VM	16MB/XA	3000	1600	QUICKDSP ON
PRODBM	1	OV/VM	16MB/XA	3000	550	QUICKDSP ON
PROMAIL	1	OV/VM	16MB/XA	3000		QUICKDSP ON
PROMBX	1	OV/VM	16MB/XA	3000		QUICKDSP ON IBCENTRL = Y
PROMBXnn	10	OV/VM	16MB/XA	3000		QUICKDSP ON IBCENTRL = Y
SMART	1	RTM	32MB/XA	3%	400	QUICKDSP ON
VMCF	1	Monitor	4MB/XA	200		QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	900	QUICKDSP ON
WRITER	1	CP Monitor	2MB/XA	100		QUICKDSP ON
Users	2100	User	2MB/XA	100		

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

Note: The OV/VM ESA Calendar Feature is not installed.

Measurement Discussion: No significant change in overall performance was observed. External response time (AVG LAST (T)) increased 1.5%, while internal throughput (ITR (H)) decreased by 0.5%. Performance improvements from the GCS pathlength and CMS working set improvements were offset by resource usage increases in other areas, particularly CP.

The CMS working set size improvement resulted in a 21-page decrease in per-user DASD page slot requirements, as indicated by TOT PAGES/USER (V)) in the results table. This decrease occurs during the execution of SYSPROF EXEC during ipl CMS. It should improve the performance of scenarios where large numbers of OfficeVision users are logging on during a short period of time.

OfficeVision uses CP's VMCF for its server communications. With VM/ESA 2.2.0, VMCF no longer has to run on the master processor (see "Performance Improvements" on page 3 and "Removal of VMCF Master Processor Dependency" on page 48 for further information). This change does not affect the measured configuration because it is not master constrained. It will, however, improve performance in high n-way OfficeVision environments that currently have a master processor constraint.

Table 8 (Page 1 of 2). OfficeVision migration from VM/ESA 2.1.0 on the 9121-480

VM/ESA Release **OfficeVision Release** **Run ID**	**2.1.0** **1.3.0** **L28V2104**	**2.2.0** **1.3.0** **L29V2101**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2100**	**2100**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.051	0.055	0.004	7.84%
NONTRIV INT	0.640	0.627	-0.013	-2.03%
TOT INT	0.601	0.589	-0.012	-2.00%
TOT INT ADJ	0.575	0.562	-0.013	-2.19%
AVG FIRST (T)	0.470	0.479	0.009	1.91%
AVG LAST (T)	0.658	0.667	0.010	1.52%
Throughput				
AVG THINK (T)	46.19	46.30	0.11	0.24%
ETR	34.01	33.91	-0.10	-0.29%
ETR (T)	35.57	35.53	-0.04	-0.10%
ETR RATIO	0.956	0.954	-0.002	-0.19%
ITR (H)	39.53	39.34	-0.19	-0.48%
ITR	18.92	18.92	0.00	-0.01%
EMUL ITR	32.87	32.93	0.06	0.18%
ITRR (H)	1.000	0.995	-0.005	-0.48%
ITRR	1.000	1.000	0.000	-0.01%
Proc. Usage				
PBT/CMD (H)	50.590	50.833	0.243	0.48%
PBT/CMD	50.609	50.379	-0.230	-0.45%
CP/CMD (H)	22.976	23.165	0.189	0.82%
CP/CMD	21.368	21.390	0.022	0.10%
EMUL/CMD (H)	27.614	27.668	0.054	0.20%
EMUL/CMD	29.241	28.989	-0.252	-0.86%
Processor Util.				
TOTAL (H)	179.93	180.61	0.68	0.38%
TOTAL	180.00	179.00	-1.00	-0.56%
UTIL/PROC (H)	89.97	90.31	0.34	0.38%
UTIL/PROC	90.00	89.50	-0.50	-0.56%
TOTAL EMUL (H)	98.22	98.31	0.09	0.09%
TOTAL EMUL	104.00	103.00	-1.00	-0.96%
MASTER TOTAL (H)	90.48	90.68	0.20	0.22%
MASTER TOTAL	90.00	90.00	0.00	0.00%
MASTER EMUL (H)	38.50	38.77	0.26	0.68%
MASTER EMUL	41.00	41.00	0.00	0.00%
TVR(H)	1.83	1.84	0.01	0.28%
TVR	1.73	1.74	0.01	0.41%
Storage				
NUCLEUS SIZE (V)	2764KB	2804KB	40KB	1.45%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	79	79	0	0.00%
PGBLPGS	55220	55072	-148	-0.27%
PGBLPGS/USER	26.3	26.2	-0.1	-0.27%
TOT PAGES/USER (V)	162	141	-21	-12.96%
FREEPGS	5864	6010	146	2.49%
FREE UTIL	1.00	0.98	-0.02	-2.43%
SHRPGS	1524	1547	23	1.51%

Table 8 (Page 2 of 2). OfficeVision migration from VM/ESA 2.1.0 on the 9121-480

VM/ESA Release **OfficeVision Release** **Run ID**	**2.1.0** **1.3.0** **L28V2104**	**2.2.0** **1.3.0** **L29V2101**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2100**	**2100**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	610	614	4	0.66%
WRITES/SEC	554	558	4	0.72%
PAGE/CMD	32.727	32.986	0.259	0.79%
PAGE IO RATE (V)	169.900	176.400	6.500	3.83%
PAGE IO/CMD (V)	4.777	4.965	0.188	3.93%
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	21.171	21.418	0.247	1.17%
Queues				
DISPATCH LIST	44.48	40.83	-3.65	-8.20%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	865	857	-8	-0.92%
VIO/CMD	24.320	24.120	-0.200	-0.82%
RIO RATE (V)	526	534	8	1.52%
RIO/CMD (V)	14.789	15.029	0.240	1.62%
NONPAGE RIO/CMD (V)	10.012	10.065	0.052	0.52%
DASD RESP TIME (V)	21.100	20.900	-0.200	-0.95%
MDC REAL SIZE (MB)	54.1	55.0	0.9	1.73%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	297	290	-7	-2.36%
MDC WRITES (I/Os)	59	59	0	0.00%
MDC AVOID	274	267	-7	-2.55%
MDC HIT RATIO	0.92	0.92	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	21.146	21.181	0.035	0.16%
DIAG/CMD	90.432	90.422	-0.010	-0.01%
DIAG 04/CMD	4.546	4.551	0.005	0.10%
DIAG 08/CMD	12.095	12.076	-0.019	-0.16%
DIAG 0C/CMD	5.419	4.737	-0.682	-12.58%
DIAG 14/CMD	1.471	1.467	-0.004	-0.28%
DIAG 58/CMD	2.060	2.060	0.000	0.01%
DIAG 98/CMD	1.294	1.295	0.001	0.05%
DIAG A4/CMD	11.132	10.954	-0.178	-1.60%
DIAG A8/CMD	6.394	6.382	-0.012	-0.19%
DIAG 214/CMD	34.052	34.243	0.191	0.56%
DIAG 270/CMD	0.000	0.647	0.647	na
SIE/CMD	130.571	130.676	0.105	0.08%
SIE INTCPT/CMD	90.094	90.166	0.072	0.08%
FREE TOTL/CMD	146.119	148.125	2.007	1.37%
VTAM Machines				
WKSET (V)	900	890	-10	-1.11%
TOT CPU/CMD (V)	5.1858	5.0973	-0.0885	-1.71%
CP CPU/CMD (V)	1.9837	1.9858	0.0021	0.11%
VIRT CPU/CMD (V)	3.2021	3.1115	-0.0906	-2.83%
DIAG 98/CMD (V)	1.294	1.290	-0.004	-0.29%

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

VSE/ESA Guest

This section examines the performance impact of migrating a VSE/ESA 2.1.0 guest from VM/ESA 2.1.0 to VM/ESA 2.2.0 All measurements were made on a 9121-320 using the DYNAPACE workload. DYNAPACE is a batch workload and is characterized by heavy I/O. See Appendix A, “Workloads” on page 70 for a description of this workload.

Measurements were obtained with the VSE/ESA* system run as a V=R guest and as a V=V guest. The V=R guest environment had dedicated DASD with I/O assist. The V=V guest environment was configured with full pack minidisk DASD with minidisk caching (MDC) active.

Workload: DYNAPACE

Hardware Configuration

Processor models: 9121-320
Storage
  Real: 256MB
  Expanded: 0MB
DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*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
SMART	1	RTM	16MB/370	100		
WRITER	1	CP monitor	2MB/XA	100		

Additional Information: 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. There is 256MB total real storage on the processor so 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 the measurement results tables.

Measurement Discussion: For both the V=R and V=V guest results, performance was equivalent to VM/ESA 2.1.0 within measurement variability.

Table 9 (Page 1 of 2). VSE/ESA V=R guest migration from VM/ESA 2.1.0 on the 9121-320

Release **Run ID**	**2.1.0** **L1R88PF3**	**2.2.0** **L1R98PF1**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96MB**	**96MB**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	877.0	862.0	-15.0	-1.71%
ETR (C)	7.66	7.80	0.13	1.74%
ITR (H)	18.96	18.94	-0.01	-0.08%
ITR	19.16	19.01	-0.14	-0.74%
ITRR (H)	1.000	0.999	-0.001	-0.08%
ITRR	1.000	0.993	-0.007	-0.74%
Proc. Usage (Sec)				
PBT/CMD (H)	3.165	3.167	0.002	0.08%
PBT/CMD	3.132	3.156	0.023	0.75%
CP/CMD (H)	0.284	0.277	-0.007	-2.41%
CP/CMD	0.235	0.231	-0.004	-1.71%
EMUL/CMD (H)	2.881	2.890	0.009	0.32%
EMUL/CMD	2.897	2.925	0.027	0.95%
Processor Util.				
TOTAL (H)	40.42	41.15	0.74	1.82%
TOTAL	40.00	41.00	1.00	2.50%
TOTAL EMUL (H)	36.79	37.55	0.76	2.07%
TOTAL EMUL	37.00	38.00	1.00	2.70%
TVR(H)	1.10	1.10	0.00	-0.24%
TVR	1.08	1.08	0.00	-0.20%
Storage				
NUCLEUS SIZE (V)	2776KB	2820KB	44KB	1.59%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	38555	38541	-14	-0.04%
FREEPGS	85	88	3	3.53%
FREE UTIL	0.59	0.57	-0.02	-2.93%
SHRPGS	3494	3495	1	0.03%
Paging				
PAGE/CMD	0.000	0.000	0.000	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	0.000	0.000	0.000	na
I/O				
VIO RATE	1.000	1.000	0.000	0.00%
VIO/CMD	7.830	7.696	-0.134	-1.71%
RIO RATE (V)	2.000	2.000	0.000	0.00%
RIO/CMD (V)	15.661	15.393	-0.268	-1.71%
DASD IO TOTAL (V)	347901	349774	1873	0.54%
DASD IO RATE (V)	414.17	416.40	2.23	0.54%
DASD IO/CMD (V)	3243.08	3204.77	-38.31	-1.18%
MDC REAL SIZE (MB)	8.2	7.9	-0.3	-3.33%
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	0.23	0.26	0.03	13.04%

Table 9 (Page 2 of 2). VSE/ESA V=R guest migration from VM/ESA 2.1.0 on the 9121-320

Release Run ID	2.1.0 L1R88PF3	2.2.0 L1R98PF1	Difference	%Difference
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96MB**	**96MB**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD (R)	10.839	10.681	-0.159	-1.46%
DIAG/CMD (R)	607.580	607.448	-0.133	-0.02%
SIE/CMD	2740.625	2709.143	-31.482	-1.15%
SIE INTCPT/CMD	2274.719	2221.497	-53.222	-2.34%
FREE TOTL/CMD	540.295	531.054	-9.241	-1.71%

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

Table 10 (Page 1 of 2). VSE/ESA V=V guest migration from VM/ESA 2.1.0 on the 9121-320

Release Run ID	2.1.0 L1V88PF6	2.2.0 L1V98PF1	Difference	%Difference
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96MB**	**96MB**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	487.0	481.0	-6.0	-1.23%
ETR (C)	13.80	13.97	0.17	1.25%
ITR (H)	14.70	14.77	0.07	0.45%
ITR	14.68	14.71	0.03	0.18%
ITRR (H)	1.000	1.004	0.004	0.45%
ITRR	1.000	1.002	0.002	0.18%
Proc. Usage (Sec)				
PBT/CMD (H)	4.081	4.063	-0.018	-0.44%
PBT/CMD	4.087	4.080	-0.007	-0.18%
CP/CMD (H)	1.134	1.133	-0.002	-0.16%
CP/CMD	1.044	1.031	-0.013	-1.23%
EMUL/CMD (H)	2.947	2.930	-0.016	-0.55%
EMUL/CMD	3.044	3.049	0.005	0.18%
Processor Util.				
TOTAL (H)	93.85	94.60	0.75	0.80%
TOTAL	94.00	95.00	1.00	1.06%
TOTAL EMUL (H)	67.76	68.23	0.47	0.69%
TOTAL EMUL	70.00	71.00	1.00	1.43%
TVR(H)	1.38	1.39	0.00	0.11%
TVR	1.34	1.34	0.00	-0.36%

Table 10 (Page 2 of 2). VSE/ESA V=V guest migration from VM/ESA 2.1.0 on the 9121-320

Release **Run ID**	**2.1.0** **L1V88PF6**	**2.2.0** **L1V98PF1**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96MB**	**96MB**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Storage				
NUCLEUS SIZE (V)	2776KB	2820KB	44KB	1.59%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	38485	38467	-18	-0.05%
FREEPGS	103	106	3	2.91%
FREE UTIL	0.63	0.61	-0.02	-2.46%
SHRPGS	1310	1436	126	9.62%
Paging				
PAGE/CMD	134.795	115.955	-18.839	-13.98%
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	278.286	279.152	0.866	0.31%
I/O				
VIO RATE	720.000	727.000	7.000	0.97%
VIO/CMD	3130.714	3122.205	-8.509	-0.27%
RIO RATE (V)	314.000	322.000	8.000	2.55%
RIO/CMD (V)	1365.339	1382.875	17.536	1.28%
DASD IO TOTAL (V)	150241	153822	3581	2.38%
DASD IO RATE (V)	313.00	320.46	7.46	2.38%
DASD IO/CMD (V)	1361.00	1376.27	15.27	1.12%
MDC REAL SIZE (MB)	111.8	111.9	0.1	0.13%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	440	445	5	1.14%
MDC WRITES (I/Os)	217	220	3	1.38%
MDC AVOID	410	412	2	0.49%
MDC HIT RATIO	0.86	0.85	-0.01	-1.16%
PRIVOPs				
PRIVOP/CMD (R)	3131.457	3117.768	-13.689	-0.44%
DIAG/CMD (R)	451.490	450.243	-1.247	-0.28%
SIE/CMD	13701.223	13665.554	-35.670	-0.26%
SIE INTCPT/CMD	11920.064	11889.032	-31.033	-0.26%
FREE TOTL/CMD	3987.313	4036.964	49.652	1.25%

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

Migration from Other VM Releases

The performance results provided in this report apply to migration from VM/ESA 2.1.0. 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 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*

220 *VM/ESA Version 2 Release 2.0 Performance Report* (this document)

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

Most 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 1.2.2 to VM/ESA 2.2.0 will tend to affect VSE guest performance, look at the VM/ESA 1.2.2 to VM/ESA 2.1.0 comparison measurements and the VM/ESA 2.1.0 to VM/ESA 2.2.0 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 45.

The comparisons listed for the CMS-intensive environment include both minidisk-only and SFS measurements. 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 42.

Table 11. 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			
VM/ESA 2.1.0	VM/ESA 2.2.0	9121-742	220			
		9672-R53	220			
		9121-480	220	220		
		9121-320			220	

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 in Table 13 on page 46.

Table 12 summarizes the ITR relationships that were observed for the CMS-intensive environment for a number of VM release-to-release transitions:

Table 12. Approximate VM relative capacity: CMS-intensive environment

Source	Target	Case	ITRR	ITRR Derivation	Notes
VM/SP 5	VM/ESA 1.5 (370) VM/ESA 2.2.0	 9221-120	0.94 0.91	R5*R13c R5*R13a*R2*R21*R22*R210	1,4,6 1,2,5-7
VM/SP 6	VM/ESA 1.5 (370) VM/ESA 2.2.0	 9221-120	1.09 1.06	R6*R13c R6*R13a*R2*R21*R22*R210	4 2,5-7
VM/ESA 1.0 (370)	VM/ESA 1.5 (370) VM/ESA 2.2.0	 9221-120 9221-170	1.02 0.99 1.06	R13c R13a*R2*R21*R22*R210 R13b*R11*R2*R21*R22*R210	 2,5-7 4-7
VM/ESA 1.5 (370)	VM/ESA 2.2.0	9221-120 9221-170	0.97 1.04	(1/R13c)*R13a*R2*R21* R22*R210 (1/R13c)*R13b*R11* R2*R21*R22*R210	2,5-7 4-7
VM/SP HPO 5	VM/ESA 2.2.0	UP MP	1.00 1.11	RHa*R2*R21*R22*R210 RHb*R1E*R11* R2*R21*R22*R210	4,6,7 3,4,6,7
VM/XA 2.0	VM/ESA 2.2.0		1.23	RX20*RX21*R1E* R11*R2*R21*R22*R210	7
VM/XA 2.1	VM/ESA 2.2.0		1.20	RX21*R1E*R11* R2*R21*R22*R210	7
VM/ESA 1.0 ESA	VM/ESA 2.2.0		1.16	R1E*R11*R2*R21*R22*R210	7
VM/ESA 1.1	VM/ESA 2.2.0		1.11	R11*R2*R21*R22*R210	7
VM/ESA 2	VM/ESA 2.2.0		1.10	R2*R21*R22*R210	7
VM/ESA 2.1	VM/ESA 2.2.0		1.09	R21*R22*R210	7
VM/ESA 2.2	VM/ESA 2.2.0		1.06	R22*R210	7
VM/ESA 2.1.0	VM/ESA 2.2.0		1.01	R210	

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 45.

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 45 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 (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. 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 1.5 (370 Feature) ITRR is based (in part) on a comparison of VM/SP 5 to VM/ESA 1.0 (370 Feature), which showed an ITRR of 0.92. This comparison was done with 16MB of real memory. However, VM/ESA 1.0 (370 Feature) supports up to 64MB of real memory (but subject to the 16MB-line constraint). When VM/SP 5 with 16MB was compared to VM/ESA 1.0 (370 Feature) with 32MB, an ITRR of 0.98 was observed. See “CMS-Intensive Migration from VM/SP Release 5” in the *VM/ESA Release 2 Performance Report* for details.

   A VM/SP HPO example: The stated VM/SP HPO 5 to VM/ESA 2.2.0 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.

5. 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 (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*.
6. 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.
7. 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 3). If these are already compiled on your system, divide the ITRR shown by 1.07.

Table 12 on page 42 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 2.2.0. From Table 12 on page 42, the estimated ITRR for migrating from VM/ESA 2.1 to VM/ESA 2.2.0 is 1.09. Therefore, the estimated overall increase in system capacity is 1.30*1.09 = 1.42.

Table 12 on page 42 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 40.

Derivation and Supporting Data

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

The derivation column in Table 12 on page 42 shows how the stated ITR ratio was calculated. For example, the ITRR of 1.06 for migrating from VM/ESA 2.2 to VM/ESA 2.2.0 was calculated by multiplying the average ITRR for migrating from VM/ESA 2.2 to VM/ESA 2.1.0 (R22) by the average ITRR for migrating from VM/ESA 2.1.0 to VM/ESA 2.2.0 (R210): 1.05*1.01 = 1.06. R22 was calculated by averaging the ITRRs for VM measurement pairs 28 through 30 (see Table 13 on page 46). Likewise, R210 was calculated by averaging the ITRRs for VM measurement pairs 31 and 32.

For the case where the source system level is VM/ESA 1.5 (370), the term “1/R13c” resolves to “1/1.02.” This takes into account the fact that VM/ESA 1.5 (370) has a somewhat higher ITR than VM/ESA 1.0 (370). This makes the ITRR smaller when migrating to VM/ESA 2.2.0 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 12 on page 42 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 12 on page 42 are based on the following pairs of measurements:

Table 13 (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)
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[3]	90	20	1.037	
26	L26E186I	L27E186J	9121-480	224MB/32MB[3]	90	16	1.026	
27	H16E0302	H17E0303	9221-170	48MB/16MB[3]	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)

Table 13 (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 2.1.0 to VM/ESA 2.2.0: FS8F0R; Report 220								
31	S48E5700	S49E5702	9121-742	1GB/1GB	90	22	1.022	
32	L28E2001	L29E2005	9121-480	256MB	90	17	1.009	
avg							1.01	(R210)

Note: The report numbers refer to the list of VM performance reports on page 40.

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 12 on page 42.

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.

[3] 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.

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.

Removal of VMCF Master Processor Dependency

Measurements were made to show the effect of moving VMCF serialization off of the master processor in VM/ESA 2.2.0. These measurements were made on a 9121-742 (4-way) processor using a workload where users repeated a CMS Pipelines statement that used the VMC Pipelines stage to cause a VMCF message to be sent from one userid to another.[4] This was measured on both VM/ESA 2.1.0 and VM/ESA 2.2.0. The results are summarized in Table 14.

Table 14. VMCF Throughput Improvements: VMCF driver on a 9121-742

VM/ESA Release Processors	2.1.0 4	2.2.0 4	Difference	%Difference
VMCF Rate (per second)	14389	26179	11790	81.94%
CPU/VMCF (msec)	0.245	0.153	-0.092	-37.55%
Avg Processor Utilization				
Total	88.3	100	11.7	13.25%
Master Processor	99.5	100	0.5	0.50%
Alternate Processors	84.5	100	15.5	18.34%

Note: All results are from VMPRF reports.

There was an 82% improvement in throughput, as measured by the rate of VMCF requests. In the VM/ESA 2.1.0 measurement, we see that the system is constrained by the master processor. The master processor is 99.5% busy, but the alternate processors are only 84.5% busy. In the VM/ESA 2.2.0 measurement, all four processors are at 100%.

For this workload, there was also an improvement in processor usage, as shown by the 37.5% decrease in CPU/VMCF. The majority of this decrease came from eliminating the processing in the dispatcher for switching to and queueing on the master processor. In addition, the decrease in task switching improved the hit ratio in the processor high speed buffer cache.

These measurements reflect an extreme case of master processor constraint. The benefit seen in customer environments would depend on the workload, the number of processors, and the VMCF rates.

[4] This method was used instead of the CP SMSG command because CP commands have certain master processor dependencies not related to VMCF.

VM/ESA on the 9672-R53

This section shows the results of measuring the FS8F0R CMS-intensive workload on a 9672-R53 and compares them to an equivalent 9121-742 measurement. Both measurements were done with identical I/O configurations and the number of users was adjusted such that average processor utilization was about 90%.

Workload: FS8F0R

Hardware Configuration

Processor model:	9121-742 or 9672-R53 (see results)
Processors used:	4 or 5 (see results)
Storage	
Real:	1024MB (default MDC)
Expanded:	1024MB (MDC BIAS 0.1[2])
Tape:	3480 (Monitor)

DASD:

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

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

Communications:

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

Software Configuration

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

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	32MB/XA	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	5700/5400	Users	3MB/XC	100		

Measurement Discussion: Response times are similar. Paging is lower in the R53 case because there are 5% fewer users. The increase in processor time per command (PBT/CMD) reflects the R53's lower per-engine processing speed.

The results show an ITR ratio (ITRR(V)) for this workload of 0.95 relative to the 9121-742. This ITRR is somewhat lower than those measured for the PD4 and HT5 workloads used for the VM/ESA platform in the Large Systems Performance Reference (LSPR). The LSPR results show ITRRs of 1.06 and 1.01 for PD4 and

HT5 respectively. Instead, the observed ITRR for FS8F0R is similar to the average ITRR of 0.93 reported for the various MVS/ESA LSPR workloads.

These ITRR results are consistent with the pattern we have seen for other workloads for migrations from 3090, 9021, or 9121 processors to 9672 and 2003 processors. HT4 and PD5 happen to run especially well on the 9672/2003 processor architecture, resulting in especially high ITRRs relative to 3090, 9021, and 9121 processors. Our experience has been that many VM/ESA workloads do not do as well as that on 9672/2003 and instead show ITRRs relative to 3090, 9021, and 9121 that are similar to those measured for the various MVS/ESA LSPR workloads. It is for this reason that we recommend use of the MVS/ESA LSPR data in addition to the VM/ESA LSPR data when sizing a migration from a 3090, 9021, or 9121 processor to a 9672 or 2003 processor.

This phenomenon only applies to migrations from 3090, 9021, or 9121 to 9672 or 2003. That is, HT4 and PD5 are good indicators of processing power relationships on VM/ESA CMS environments for processor migrations within 3090/9021/9121 and within 9672/2003.

Hardware instrumentation was not available on the measured 9672 processor so the ITR calculations are instead based on the average processor utilization reported by VMPRF. Results on processors that do have instrumentation consistently show a very close correspondence between total processor utilization reported by hardware instrumentation, CP monitor (as reported by VMPRF), and RTM VM/ESA.

Table 15 (Page 1 of 3). Comparison of 9672-R53 to 9121-742, FS8F0R workload

Processor **VM/ESA Release** **Run ID**	**9121-742** **2.2.0** **S49E5702**	**9672-R53** **2.2.0** **C59E5400**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5700**	**5400**	**-300**	**-5.26%**
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**5**	**1**	**25.00%**
Response Time				
TRIV INT	0.105	0.095	-0.010	-9.52%
NONTRIV INT	0.350	0.342	-0.008	-2.29%
TOT INT	0.250	0.236	-0.014	-5.60%
TOT INT ADJ	0.247	0.242	-0.005	-1.84%
AVG FIRST (T)	0.266	0.286	0.021	7.78%
AVG LAST (T)	0.353	0.378	0.025	6.98%
Throughput				
AVG THINK (T)	26.13	26.14	0.01	0.04%
ETR	196.89	194.21	-2.68	-1.36%
ETR (T)	199.46	189.21	-10.25	-5.14%
ETR RATIO	0.987	1.026	0.039	3.98%
ITR (V)	225.89	213.80	-12.10	-5.36%
ITR	55.81	43.85	-11.96	-21.42%
EMUL ITR	87.49	67.58	-19.91	-22.76%
ITRR (V)	1.000	0.946	-0.054	-5.36%
ITRR	1.000	0.786	-0.214	-21.42%
Proc. Usage				
PBT/CMD	17.698	23.466	5.769	32.60%
CP/CMD	6.417	8.245	1.828	28.48%
EMUL/CMD	11.280	15.221	3.941	34.94%

Table 15 (Page 2 of 3). Comparison of 9672-R53 to 9121-742, FS8F0R workload

Processor **VM/ESA Release** **Run ID**	**9121-742** **2.2.0** **S49E5702**	**9672-R53** **2.2.0** **C59E5400**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5700**	**5400**	**-300**	**-5.26%**
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**5**	**1**	**25.00%**
Processor Util.				
TOTAL	353.00	444.00	91.00	25.78%
UTIL/PROC	88.25	88.80	0.55	0.62%
UTIL/PROC (V)	88.30	88.50	0.20	0.23%
TOTAL EMUL	225.00	288.00	63.00	28.00%
TVR	1.57	1.54	-0.03	-1.74%
Storage				
NUCLEUS SIZE (V)	2808KB	2808KB	0KB	0.00%
TRACE TABLE (V)	650KB	800KB	150KB	23.08%
WKSET (V)	75	71	-4	-5.33%
PGBLPGS	231K	231K	0K	0.00%
PGBLPGS/USER	40.5	42.8	2.3	5.56%
TOT PAGES/USER (V)	189	185	-4	-2.12%
FREEPGS	16953	16060	-893	-5.27%
FREE UTIL	0.92	0.92	0.00	0.37%
SHRPGS	1973	1885	-88	-4.46%
Paging				
READS/SEC	952	639	-313	-32.88%
WRITES/SEC	669	403	-266	-39.76%
PAGE/CMD	8.127	5.507	-2.620	-32.24%
PAGE IO RATE (V)	254.700	152.600	-102.100	-40.09%
PAGE IO/CMD (V)	1.277	0.807	-0.470	-36.84%
XSTOR IN/SEC	681	834	153	22.47%
XSTOR OUT/SEC	1516	1362	-154	-10.16%
XSTOR/CMD	11.015	11.606	0.592	5.37%
FAST CLR/CMD	8.794	8.773	-0.020	-0.23%
Queues				
DISPATCH LIST	108.55	99.30	-9.26	-8.53%
ELIGIBLE LIST	0.00	0.02	0.02	na
I/O				
VIO RATE	1815	1721	-94	-5.18%
VIO/CMD	9.099	9.096	-0.004	-0.04%
RIO RATE (V)	639	518	-121	-18.94%
RIO/CMD (V)	3.204	2.738	-0.466	-14.54%
NONPAGE RIO/CMD (V)	1.927	1.931	0.005	0.23%
DASD RESP TIME (V)	20.400	20.000	-0.400	-1.96%
MDC REAL SIZE (MB)	32.0	34.3	2.3	7.16%
MDC XSTOR SIZE (MB)	63.0	63.2	0.1	0.19%
MDC READS (I/Os)	570	540	-30	-5.26%
MDC WRITES (I/Os)	27	26	-1	-3.70%
MDC AVOID	531	503	-28	-5.27%
MDC HIT RATIO	0.93	0.93	0.00	0.00%

Table 15 (Page 3 of 3). Comparison of 9672-R53 to 9121-742, FS8F0R workload

Processor	**9121-742**	**9672-R53**		
VM/ESA Release	**2.2.0**	**2.2.0**	**Difference**	**%Difference**
Run ID	**S49E5702**	**C59E5400**		
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5700**	**5400**	**-300**	**-5.26%**
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**5**	**1**	**25.00%**
PRIVOPs				
PRIVOP/CMD	20.479	20.135	-0.344	-1.68%
DIAG/CMD	24.430	24.407	-0.023	-0.09%
DIAG 04/CMD	0.904	0.896	-0.008	-0.85%
DIAG 08/CMD	0.736	0.736	0.001	0.07%
DIAG 0C/CMD	0.212	0.212	0.000	-0.03%
DIAG 14/CMD	0.025	0.025	0.000	0.05%
DIAG 58/CMD	1.249	1.249	0.000	0.02%
DIAG 98/CMD	0.298	0.294	-0.005	-1.56%
DIAG A4/CMD	3.559	3.553	-0.006	-0.18%
DIAG A8/CMD	2.664	2.668	0.004	0.14%
DIAG 214/CMD	12.485	12.479	-0.006	-0.05%
DIAG 270/CMD	0.921	0.921	0.000	0.00%
SIE/CMD	55.148	52.545	-2.603	-4.72%
SIE INTCPT/CMD	36.949	36.781	-0.168	-0.45%
FREE TOTL/CMD	44.891	44.897	0.007	0.02%
VTAM Machines				
WKSET (V)	4140	4118	-22	-0.53%
TOT CPU/CMD (V)	2.8688	4.2369	1.3681	47.69%
CP CPU/CMD (V)	1.2617	1.9115	0.6498	51.50%
VIRT CPU/CMD (V)	1.6071	2.3255	0.7184	44.70%
DIAG 98/CMD (V)	0.297	0.294	-0.003	-1.04%

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

VM/ESA on the PC Server 520 and the RS/6000 591

A set of FS8F CMS-intensive measurements was obtained on the PC Server 520 and the RS/6000 591 with 128MB of S/390 memory. A report, *VM/ESA Performance on the P/390 and R/390* (see "Referenced Publications"), has been prepared that presents the results of these measurements, compares them to similar measurements taken in 1995 on the PC Server 500, and makes a number of tuning recommendations. This section summarizes the main findings from that report.

The S/390 Microprocessor Complex is identical in the PC Server 500, the PC Server 520, and the RS/6000 591. In addition, the same release of VM/ESA (1.2.2) was used for all the measurements. Therefore, all observed performance differences are due to differences in the I/O characteristics as determined by improvements to the P/390 support code, the workstation hardware, software, adapters, and I/O configuration. These differences are listed in the report. Each configuration was built using recommended components that were generally available at the time the measurements were done.

PC Server 520: The 1995 measurements showed that the PC Server 500 could run 190 FS8F CMS users with an average external response time of 1.02 seconds. The measured PC Server 520 configuration ran the same 190 users with an average response time of 0.81 seconds (a 20% improvement) when both cases used a 2MB OS/2 cache — the maximum supported by HPFS.

The Server 520 configuration used OS/2 Warp Server Advanced, which includes the HPFS386 file system. HPFS386 supports OS/2 cache sizes greater than 2MB. When a cache size of 17MB was used, average response time decreased to 0.73 seconds (a 28% improvement).

This improved response time allowed the number of users to be increased somewhat. A 200 user measurement resulted in a 0.90 second average response time. At that loading, the S/390 utilization was 84%, indicating limited potential for increasing the capacity further.

RS/6000 591: The measured RS/6000 591 configuration ran 190 FS8F CMS users with an average response time of 0.51 seconds. This represents a 30% improvement relative to the corresponding PC Server 520 run with 17MB of HPFS386 cache and a 50% improvement relative to the 190 user measurement on the PC Server 500.

Although the RS/6000 591 offered better response time, S/390 capacity was about the same. It supported 200 FS8F CMS users with an average response time of 0.69. However, 220 users ran with an average response time of 1.56 seconds -- well above the subsecond range. Capacity was similar because it was determined by the speed of the S/390 card, which is the same on all three systems.

Configuration of Storage as Central or Expanded

Recent S/390 processor families (such as 9672) provide flexibility as to how much of the processor's electronic storage is used as central storage and how much is used as expanded storage. How this storage is configured will affect the performance of a VM/ESA system running on that processor.

Two CMS-intensive measurements were obtained on a 2GB 9672-R53 to quantify the performance differences between configuring this system with 1GB central / 1GB expanded and configuring it with 2GB central and no expanded storage.

Workload: FS8F0R

Hardware Configuration

Processor model:	9672-R53
Processors used:	5
Storage	
Real:	1024MB or 2048MB (see results); default MDC
Expanded:	1024MB[5] or 0MB (see results)
Tape:	3480 (Monitor)

DASD:

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

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

Communications:

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

Software Configuration

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

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
SMART	1	RTM	32MB/XA	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	5400	Users	3MB/XC	100		

[5] The expanded storage MDC BIAS was set to 0.1 (using the SET MDCACHE command). Previous measurements have shown that this improves overall system performance for this particular environment. For more information refer to the *VM/ESA Release 2.2 Performance Report*.

Measurement Discussion: Overall performance is quite similar, with average external (end user) response time (AVG LAST (T)) being slightly higher (1%) and processor usage per command (PBT/CMD) being slightly lower (-0.4%) for the no expanded storage case.

The slight decrease in PBT/CMD reflects the fact that CP is somewhat more efficient when there is no expanded storage because it does not have to move pages to/from expanded storage and there is no page migration.

The much larger internal response time increases reported by RTM (such as TRIV INT and TOT INT) and also by VMPRF do not, for the most part, translate into corresponding increases in end user response time.

Part of the reason why the percentage increases are higher for internal response times is that the storage change has resulted in a shift in what the CP scheduler considers to be a transaction. This is shown by the 3.4% throughput (ETR) decrease reported by RTM. (The true throughput rate (ETR(T)) measured by TPNS actually increased by 1.3%.) Average internal response time after adjustment for this shift (TOT INT ADJ) shows a 14% increase, down from 20% for the corresponding unadjusted figure (TOT INT).

The remaining “true” increase in internal response time apparently has been largely offset by a reduction in the portion of external response time that is not counted by internal response time. This includes time spent in VTAM/VSCS and network latency.

Average working set (WKSET(V)) and average central storage + expanded storage + DASD page slots per user (TOT PAGES/USER(V)) are both higher in the no expanded storage case because all paging is being done to DASD, which makes use of block paging. Block paging increases working sets somewhat because some of the pages that are paged in as part of the block turn out to be no longer needed. Block paging is not used for paging to/from expanded storage.

Hardware instrumentation was not available on the measured 9672 processor.

Table 16 (Page 1 of 2). Effect of storage configuration on 9672-R53 CMS-intensive performance

Release **Run ID**	**2.2.0** **C59E5400**	**2.2.0** **C59E5401**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**2048MB**		
Exp. Storage	**1024MB**	**0MB**		
Users	**5400**	**5400**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**5**	**5**		
Response Time				
TRIV INT	0.095	0.125	0.030	31.58%
NONTRIV INT	0.342	0.391	0.049	14.33%
TOT INT	0.236	0.283	0.047	19.92%
TOT INT ADJ	0.242	0.277	0.035	14.35%
AVG FIRST (T)	0.286	0.292	0.006	1.98%
AVG LAST (T)	0.378	0.382	0.004	0.97%
Throughput				
AVG THINK (T)	26.14	25.87	-0.27	-1.05%
ETR	194.21	187.55	-6.66	-3.43%
ETR (T)	189.21	191.61	2.40	1.27%
ETR RATIO	1.026	0.979	-0.048	-4.64%
ITR (V)	213.80	214.09	0.30	0.14%
ITR	43.85	41.89	-1.96	-4.46%
EMUL ITR	67.58	64.55	-3.03	-4.48%
ITRR (V)	1.000	1.001	0.001	0.14%
ITRR	1.000	0.955	-0.045	-4.46%
Proc. Usage				
PBT/CMD	23.466	23.381	-0.086	-0.36%
CP/CMD	8.245	8.194	-0.051	-0.62%
EMUL/CMD	15.221	15.187	-0.034	-0.23%
Processor Util.				
TOTAL	444.00	448.00	4.00	0.90%
UTIL/PROC	88.80	89.60	0.80	0.90%
UTIL/PROC (V)	88.50	89.50	1.00	1.13%
TOTAL EMUL	288.00	291.00	3.00	1.04%
TVR	1.54	1.54	0.00	-0.14%
Storage				
NUCLEUS SIZE (V)	2808KB	2808KB	0KB	0.00%
TRACE TABLE (V)	800KB	800KB	0KB	0.00%
WKSET (V)	71	89	18	25.35%
PGBLPGS	231K	492K	261K	112.99%
PGBLPGS/USER	42.8	91.1	48.3	112.99%
TOT PAGES/USER (V)	185	207	22	11.89%
FREEPGS	16060	16127	67	0.42%
FREE UTIL	0.92	0.92	0.00	-0.42%
SHRPGS	1885	1931	46	2.44%
Paging				
READS/SEC	639	1413	774	121.13%
WRITES/SEC	403	1013	610	151.36%
PAGE/CMD	5.507	12.661	7.154	129.90%
PAGE IO RATE (V)	152.600	409.800	257.200	168.55%
PAGE IO/CMD (V)	0.807	2.139	1.332	165.18%
XSTOR IN/SEC	834	0	-834	-100.00%
XSTOR OUT/SEC	1362	0	-1362	-100.00%
XSTOR/CMD	11.606	0.000	-11.606	-100.00%
FAST CLR/CMD	8.773	8.768	-0.006	-0.06%
Queues				
DISPATCH LIST	99.30	114.44	15.15	15.26%
ELIGIBLE LIST	0.02	0.00	-0.02	-100.00%

Table 16 (Page 2 of 2). Effect of storage configuration on 9672-R53 CMS-intensive performance

Release **Run ID**	**2.2.0** **C59E5400**	**2.2.0** **C59E5401**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**2048MB**		
Exp. Storage	**1024MB**	**0MB**		
Users	**5400**	**5400**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**5**	**5**		
I/O				
VIO RATE	1721	1742	21	1.22%
VIO/CMD	9.096	9.091	-0.004	-0.05%
RIO RATE (V)	518	767	249	48.07%
RIO/CMD (V)	2.738	4.003	1.265	46.21%
NONPAGE RIO/CMD (V)	1.931	1.864	-0.067	-3.47%
DASD RESP TIME (V)	20.000	21.200	1.200	6.00%
MDC REAL SIZE (MB)	34.3	371.4	337.1	983.72%
MDC XSTOR SIZE (MB)	63.2	0.0	-63.2	-100.00%
MDC READS (I/Os)	540	547	7	1.30%
MDC WRITES (I/Os)	26	37	11	42.31%
MDC AVOID	503	519	16	3.18%
MDC HIT RATIO	0.93	0.94	0.01	1.08%
PRIVOPs				
PRIVOP/CMD	20.135	20.139	0.004	0.02%
DIAG/CMD	24.407	24.436	0.029	0.12%
DIAG 04/CMD	0.896	0.883	-0.013	-1.46%
DIAG 08/CMD	0.736	0.737	0.000	0.04%
DIAG 0C/CMD	0.212	0.212	0.000	0.15%
DIAG 14/CMD	0.025	0.025	0.000	0.16%
DIAG 58/CMD	1.249	1.250	0.001	0.08%
DIAG 98/CMD	0.294	0.278	-0.015	-5.27%
DIAG A4/CMD	3.553	3.556	0.004	0.10%
DIAG A8/CMD	2.668	2.679	0.010	0.39%
DIAG 214/CMD	12.479	12.520	0.041	0.33%
DIAG 270/CMD	0.921	0.921	0.000	-0.02%
SIE/CMD	52.545	49.428	-3.117	-5.93%
SIE INTCPT/CMD	36.781	37.071	0.289	0.79%
FREE TOTL/CMD	44.897	44.199	-0.699	-1.56%
VTAM Machines				
WKSET (V)	4118	4125	7	0.17%
TOT CPU/CMD (V)	4.2369	4.2766	0.0397	0.94%
CP CPU/CMD (V)	1.9115	1.9339	0.0224	1.17%
VIRT CPU/CMD (V)	2.3255	2.3427	0.0172	0.74%
DIAG 98/CMD (V)	0.294	0.278	-0.016	-5.30%

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

TCP/IP 2.4.0

This section summarizes the results from a pair of measurements obtained to observe the effects of migrating from TCP/IP 2.3.0 to TCP/IP 2.4.0. These measurements were made using VM/ESA 2.1.0 running on a 9121-320.

RPC throughput measurements were used for this evaluation. These measurements are equivalent to the 4-thread, 9121-320 measurements described in the *VM/ESA Version 2 Release 1.0 Performance Report* ("OpenEdition DCE for VM/ESA" chapter). This workload exercises the mainline path through UDP and IP in the TCP/IP protocol stack.

The measurement configuration consisted of an AIX* DCE client running on an RS/6000 and a VM/ESA DCE server running on a 9121-320. These systems were connected by a 16 megabit IBM Token Ring with the 9121-320 connected to the token ring through a 3172-1 control unit. The client application issued 1KB null RPC requests with zero think time between requests from each of four separate threads. This created a load that was sufficient to fully utilize the 9121-320 server configuration. The TCPIP and DCE server virtual machines accounted for nearly all of the processor usage.

CP monitor records were collected on the VM/ESA server system and later reduced using VMPRF. Total and virtual CPU-seconds used by the TCPIP virtual machine during the measurement were obtained from the USER_RESOURCE_UTIL report. A total messages count (IUCV sends + IUCV receives) was obtained from the VMCOMM_ACTIVITY_BY_TIME report. The results are summarized in Table 17.

Table 17. TCP/IP CPU usage: 1KB RPCs to a 9121-320 server

Run ID **TCP/IP Release**	**TCP23B** **2.3.0**	**TCP24A** **2.4.0**
Run duration (seconds)	300	300
TCP/IP CPU usage (seconds)		
Total	98	93
Virtual	38	36
Total messages	348290	335983
TCP/IP CPU/message (msec)		
Total	0.28	0.28
Virtual	0.11	0.11
CP	0.17	0.17

Note: All results are from VMPRF reports.

The results show no measurable difference in TCP/IP processor usage between TCP/IP 2.3.0 and TCP/IP 2.4.0.

OfficeVision 1.3.0

This section summarizes the results from a pair of measurements obtained to observe the effects of migrating from OfficeVision 1.2.0 to OfficeVision 1.3.0 for the IBM Office Benchmark (IOB). Both measurements were made using VM/ESA 2.1.0 running on a 9121-480.

Workload: IOB

Hardware Configuration

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

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes - PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	6	4	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4	10	4	2	16 R	7 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: IOB
CMS block size: 4KB

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
Ennnn	10	Workload	2MB/XA	1000		QUICKDSP ON
PRNTEAT1	1	Workload	2MB/XA	1000		QUICKDSP ON
PROCAL	1	OV/VM	16MB/XA	3000	1600	QUICKDSP ON
PRODBM	1	OV/VM	16MB/XA	3000	550	QUICKDSP ON
PROMAIL	1	OV/VM	16MB/XA	3000		QUICKDSP ON
PROMBX	1	OV/VM	16MB/XA	3000		QUICKDSP ON IBCENTRL = Y
PROMBXnn	10	OV/VM	16MB/XA	3000		QUICKDSP ON IBCENTRL = Y
SMART	1	RTM	32MB/XA	3%	400	QUICKDSP ON
VMCF	1	Monitor	4MB/XA	200		QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	900	QUICKDSP ON
WRITER	1	CP Monitor	2MB/XA	100		QUICKDSP ON
Users	2100	User	2MB/XA	100		

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

Note: The OV/VM ESA Calendar Feature is not installed.

Measurement Discussion: The results for the two OfficeVision releases are equivalent within measurement variability.

Table 18 (Page 1 of 2). Migration from OfficeVision 1.2.0 to OfficeVision 1.3.0 using VM/ESA 2.1.0 on a 9121-480

OV/VM Release **VM/ESA Release** **Run ID**	**1.2.0** **2.1.0** **L28V2103**	**1.3.0** **2.1.0** **L28V2104**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2100**	**2100**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.050	0.051	0.001	2.00%
NONTRIV INT	0.639	0.640	0.001	0.16%
TOT INT	0.600	0.601	0.001	0.17%
TOT INT ADJ	0.573	0.575	0.002	0.32%
AVG FIRST (T)	0.479	0.470	-0.009	-1.77%
AVG LAST (T)	0.661	0.658	-0.004	-0.53%
Throughput				
AVG THINK (T)	46.39	46.19	-0.20	-0.43%
ETR	33.93	34.01	0.08	0.24%
ETR (T)	35.54	35.57	0.03	0.08%
ETR RATIO	0.955	0.956	0.001	0.15%
ITR (H)	39.43	39.53	0.11	0.27%
ITR	18.84	18.92	0.08	0.42%
EMUL ITR	32.73	32.87	0.14	0.42%
ITRR (H)	1.000	1.003	0.003	0.27%
ITRR	1.000	1.004	0.004	0.42%
Proc. Usage				
PBT/CMD (H)	50.725	50.590	-0.135	-0.27%
PBT/CMD	50.651	50.609	-0.042	-0.08%
CP/CMD (H)	23.029	22.976	-0.053	-0.23%
CP/CMD	21.386	21.368	-0.018	-0.08%
EMUL/CMD (H)	27.696	27.614	-0.081	-0.29%
EMUL/CMD	29.265	29.241	-0.024	-0.08%
Processor Util.				
TOTAL (H)	180.26	179.93	-0.33	-0.18%
TOTAL	180.00	180.00	0.00	0.00%
UTIL/PROC (H)	90.13	89.97	-0.17	-0.18%
UTIL/PROC	90.00	90.00	0.00	0.00%
TOTAL EMUL (H)	98.42	98.22	-0.21	-0.21%
TOTAL EMUL	104.00	104.00	0.00	0.00%
MASTER TOTAL (H)	90.59	90.48	-0.10	-0.11%
MASTER TOTAL	91.00	90.00	-1.00	-1.10%
MASTER EMUL (H)	38.60	38.50	-0.09	-0.24%
MASTER EMUL	41.00	41.00	0.00	0.00%
TVR(H)	1.83	1.83	0.00	0.03%
TVR	1.73	1.73	0.00	0.00%
Storage				
NUCLEUS SIZE (V)	2764KB	2764KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	80	79	-1	-1.25%
PGBLPGS	55228	55220	-8	-0.01%
PGBLPGS/USER	26.3	26.3	0.0	-0.01%
FREEPGS	5855	5864	9	0.15%
FREE UTIL	1.01	1.00	0.00	-0.15%
SHRPGS	1512	1524	12	0.79%

Table 18 (Page 2 of 2). Migration from OfficeVision 1.2.0 to OfficeVision 1.3.0 using VM/ESA 2.1.0 on a 9121-480

OV/VM Release **VM/ESA Release** **Run ID**	**1.2.0** **2.1.0** **L28V2103**	**1.3.0** **2.1.0** **L28V2104**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2100**	**2100**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	620	610	-10	-1.61%
WRITES/SEC	576	554	-22	-3.82%
PAGE/CMD	33.655	32.727	-0.928	-2.76%
PAGE IO RATE (V)	162.500	169.900	7.400	4.55%
PAGE IO/CMD (V)	4.573	4.777	0.204	4.47%
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	21.442	21.171	-0.271	-1.26%
Queues				
DISPATCH LIST	44.83	44.48	-0.35	-0.78%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	876	865	-11	-1.26%
VIO/CMD	24.650	24.320	-0.330	-1.34%
RIO RATE (V)	519	526	7	1.35%
RIO/CMD (V)	14.604	14.789	0.185	1.26%
NONPAGE RIO/CMD (V)	10.032	10.012	-0.020	-0.19%
DASD RESP TIME (V)	21.300	21.100	-0.200	-0.94%
MDC REAL SIZE (MB)	54.8	54.1	-0.8	-1.41%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	309	297	-12	-3.88%
MDC WRITES (I/Os)	58	59	1	1.72%
MDC AVOID	286	274	-12	-4.20%
MDC HIT RATIO	0.92	0.92	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	21.161	21.146	-0.015	-0.07%
DIAG/CMD	90.122	90.432	0.310	0.34%
DIAG 04/CMD	4.466	4.546	0.080	1.80%
DIAG 08/CMD	12.128	12.095	-0.032	-0.27%
DIAG 0C/CMD	5.428	5.419	-0.008	-0.16%
DIAG 14/CMD	1.474	1.471	-0.003	-0.21%
DIAG 58/CMD	2.060	2.060	-0.001	-0.03%
DIAG 98/CMD	1.282	1.294	0.012	0.93%
DIAG A4/CMD	11.492	11.132	-0.359	-3.13%
DIAG A8/CMD	6.379	6.394	0.015	0.23%
DIAG 214/CMD	33.923	34.052	0.128	0.38%
SIE/CMD	130.200	130.571	0.370	0.28%
SIE INTCPT/CMD	89.838	90.094	0.255	0.28%
FREE TOTL/CMD	145.564	146.119	0.554	0.38%
VTAM Machines				
WKSET (V)	900	900	0	0.00%
TOT CPU/CMD (V)	5.1901	5.1858	-0.0043	-0.08%
CP CPU/CMD (V)	1.9854	1.9837	-0.0017	-0.09%
VIRT CPU/CMD (V)	3.2047	3.2021	-0.0026	-0.08%
DIAG 98/CMD (V)	1.282	1.294	0.012	0.90%

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

OfficeVision SFS A-Directory Support

OfficeVision 1.4.0 supports the use of an SFS directory as filemode A. One of the primary benefits of this support is that it enables OfficeVision/VM installations to significantly reduce their DASD space requirements by moving user files from minidisks to SFS. DASD space savings exceeding 30% have been reported in the past for non-OfficeVision CMS systems. These improvements and all other SFS benefits need to be traded off against the fact that SFS uses additional processor time and real storage relative to minidisks.

This section quantifies the performance effects of replacing the A-minidisk with an SFS A-directory for the case of the IOB workload running on a 9121-480 processor. It then provides a generalized method for estimating the percentage increase in processor usage for moving a given amount of minidisk activity to SFS filecontrol directories.

Workload: IOB

Hardware Configuration

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

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes -* *PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-2	3990-2	4	6	4	6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4	10	4	2	16 R	7 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: IOB
CMS block size: 4KB

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
Ennnn	10	Workload	2MB/XA	1000		QUICKDSP ON
PRNTEAT1	1	Workload	2MB/XA	1000		QUICKDSP ON
PROCAL	1	OV/VM	16MB/XA	3000	1600	QUICKDSP ON
PRODBM	1	OV/VM	16MB/XA	3000	550	QUICKDSP ON
PROMAIL	1	OV/VM	16MB/XA	3000		QUICKDSP ON
PROMBX	1	OV/VM	16MB/XA	3000		QUICKDSP ON IBCENTRL = Y
PROMBXnn	10	OV/VM	16MB/XA	3000		QUICKDSP ON IBCENTRL = Y
CRRSERV1	1	SFS	16MB/XC	100		
ROSERV1	1	SFS	64MB/XC	100		QUICKDSP ON
RWSERVn	2	SFS	64MB/XC	1500	1300	QUICKDSP ON
SMART	1	RTM	32MB/XA	3%	400	QUICKDSP ON
VMCF	1	Monitor	4MB/XA	200		QUICKDSP ON
VTAMXA	1	VTAM/VSCS	64MB/XA	10000	900	QUICKDSP ON
WRITER	1	CP Monitor	2MB/XA	100		QUICKDSP ON
Users	2100/1500	User	2MB/XA	100		

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

Note: The OV/VM ESA Calendar Feature is not installed.

Measurement Discussion: An OfficeVision measurement was obtained with each user's SFS root directory used for that user's A-filemode. The number of users (1500) was chosen such that average processor utilization was approximately 90%. This measurement is otherwise equivalent to the 2100 user VM/ESA 2.2.0 run (L29V2101) shown in Table 8 on page 34. These two measurements are compared in Table 19. Both measurements were made using OfficeVision 1.3.0.[6]

[6] OfficeVision 1.3.0 does not support an SFS directory as filemode A but the FS8F workload does run correctly in that environment. Similar comparative results can be expected using OfficeVision 1.4.0.

Table 19 (Page 1 of 3). OfficeVision performance with an SFS A-directory

Filemode A	**Minidisk**	**SFS**	**Difference**	**%Difference**
OV/VM Release	**1.3.0**	**1.3.0**		
VM/ESA Release	**2.2.0**	**2.2.0**		
Run ID	**L29V2101**	**L29V1500**		
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2100**	**1500**	**-600**	**-28.57%**
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.055	0.121	0.066	120.00%
NONTRIV INT	0.627	0.728	0.101	16.11%
TOT INT	0.589	0.683	0.094	15.96%
TOT INT ADJ	0.562	0.660	0.098	17.39%
AVG FIRST (T)	0.479	0.525	0.045	9.48%
AVG LAST (T)	0.667	0.759	0.092	13.78%
Throughput				
AVG THINK (T)	46.30	46.39	0.09	0.18%
ETR	33.91	24.55	-9.36	-27.60%
ETR (T)	35.53	25.41	-10.12	-28.48%
ETR RATIO	0.954	0.966	0.012	1.23%
ITR (H)	39.34	28.45	-10.89	-27.68%
ITR	18.92	13.75	-5.16	-27.30%
EMUL ITR	32.93	24.17	-8.76	-26.60%
ITRR (H)	1.000	0.723	-0.277	-27.68%
ITRR	1.000	0.727	-0.273	-27.30%
Proc. Usage				
PBT/CMD (H)	50.833	70.291	19.457	38.28%
PBT/CMD	50.379	70.442	20.064	39.83%
CP/CMD (H)	23.165	32.129	8.964	38.70%
CP/CMD	21.390	30.302	8.912	41.67%
EMUL/CMD (H)	27.668	38.162	10.493	37.92%
EMUL/CMD	28.989	40.140	11.151	38.47%
Processor Util.				
TOTAL (H)	180.61	178.61	-2.00	-1.11%
TOTAL	179.00	179.00	0.00	0.00%
UTIL/PROC (H)	90.31	89.31	-1.00	-1.11%
UTIL/PROC	89.50	89.50	0.00	0.00%
TOTAL EMUL (H)	98.31	96.97	-1.34	-1.36%
TOTAL EMUL	103.00	102.00	-1.00	-0.97%
MASTER TOTAL (H)	90.68	89.58	-1.10	-1.21%
MASTER TOTAL	90.00	90.00	0.00	0.00%
MASTER EMUL (H)	38.77	41.74	2.97	7.67%
MASTER EMUL	41.00	44.00	3.00	7.32%
TVR(H)	1.84	1.84	0.00	0.26%
TVR	1.74	1.75	0.02	0.98%
Storage				
NUCLEUS SIZE (V)	2804KB	2804KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	79	92	13	16.46%
PGBLPGS	55072	57221	2149	3.90%
PGBLPGS/USER	26.2	38.1	11.9	45.46%
TOT PAGES/USER (V)	141	171	30	21.28%
FREEPGS	6010	4398	-1612	-26.82%
FREE UTIL	0.98	0.99	0.01	1.00%
SHRPGS	1547	1707	160	10.34%

Table 19 (Page 2 of 3). OfficeVision performance with an SFS A-directory

Filemode A	**Minidisk**	**SFS**	**Difference**	**%Difference**
OV/VM Release	**1.3.0**	**1.3.0**		
VM/ESA Release	**2.2.0**	**2.2.0**		
Run ID	**L29V2101**	**L29V1500**		
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2100**	**1500**	**-600**	**-28.57%**
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	614	504	-110	-17.92%
WRITES/SEC	558	456	-102	-18.28%
PAGE/CMD	32.986	37.779	4.794	14.53%
PAGE IO RATE (V)	176.400	143.600	-32.800	-18.59%
PAGE IO/CMD (V)	4.965	5.651	0.686	13.83%
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	21.418	21.329	-0.089	-0.41%
Queues				
DISPATCH LIST	40.83	35.74	-5.10	-12.48%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	857	676	-181	-21.12%
VIO/CMD	24.120	26.603	2.483	10.29%
RIO RATE (V)	534	433	-101	-18.91%
RIO/CMD (V)	15.029	17.040	2.011	13.38%
NONPAGE RIO/CMD (V)	10.065	11.389	1.324	13.16%
DASD RESP TIME (V)	20.900	17.700	-3.200	-15.31%
MDC REAL SIZE (MB)	55.0	68.0	13.0	23.69%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	290	244	-46	-15.86%
MDC WRITES (I/Os)	59	52	-7	-11.86%
MDC AVOID	267	215	-52	-19.48%
MDC HIT RATIO	0.92	0.87	-0.05	-5.43%
PRIVOPs				
PRIVOP/CMD	21.181	50.818	29.638	139.93%
DIAG/CMD	90.422	85.133	-5.289	-5.85%
DIAG 04/CMD	4.551	5.515	0.964	21.18%
DIAG 08/CMD	12.076	12.822	0.746	6.18%
DIAG 0C/CMD	4.737	5.072	0.334	7.06%
DIAG 14/CMD	1.467	1.636	0.169	11.54%
DIAG 58/CMD	2.060	2.111	0.052	2.50%
DIAG 98/CMD	1.295	1.619	0.324	25.05%
DIAG A4/CMD	10.954	6.018	-4.936	-45.06%
DIAG A8/CMD	6.382	5.406	-0.976	-15.29%
DIAG 214/CMD	34.243	32.050	-2.193	-6.40%
DIAG 270/CMD	0.673	0.685	0.012	1.79%
SIE/CMD	130.676	168.472	37.796	28.92%
SIE INTCPT/CMD	90.166	124.669	34.503	38.27%
FREE TOTL/CMD	148.125	164.182	16.057	10.84%
VTAM Machines				
WKSET (V)	890	900	10	1.12%
TOT CPU/CMD (V)	5.0973	5.5313	0.4340	8.51%
CP CPU/CMD (V)	1.9858	2.1426	0.1568	7.90%
VIRT CPU/CMD (V)	3.1115	3.3888	0.2773	8.91%
DIAG 98/CMD (V)	1.290	1.619	0.329	25.50%

Table 19 (Page 3 of 3). OfficeVision performance with an SFS A-directory

Filemode A	**Minidisk**	**SFS**		
OV/VM Release	**1.3.0**	**1.3.0**	**Difference**	**%Difference**
VM/ESA Release	**2.2.0**	**2.2.0**		
Run ID	**L29V2101**	**L29V1500**		
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2100**	**1500**	**-600**	**-28.57%**
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
SFS Servers				
WKSET (V)	na	2623		
TOT CPU/CMD (V)	na	14.514		
CP CPU/CMD (V)	na	6.141		
VIRT CPU/CMD (V)	na	8.373		
FP REQ/CMD(Q)	na	5.577		
IO/CMD (Q)	na	6.683		
IO TIME/CMD (Q)	na	0.082		
SFS TIME/CMD (Q)	na	0.153		

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

The results show the same kinds of performance effects that have been observed for the FS8F CMS-intensive workload. See Table 7 on page 29 for an example FS8F comparison. Processor usage per command (PBT/CMD(H)) increased by 38%. This is higher than the 19% increase observed for the FS8F workload. The reason for this is that a correspondingly larger amount of minidisk activity has been moved to SFS filecontrol directories in the IOB workload case.

A rule of thumb has been developed, based on earlier FS8F workload results, for estimating the percentage increase in processor usage for migrating a given amount of minidisk activity to SFS filecontrol directories:

> Assume a 6% increase in total system CPU usage for every virtual I/O per million instructions executed that is moved from minidisks to SFS filecontrol directories.

A VM/ESA installation can use this method to estimate the effect of their own planned usage of SFS on their system's total processor requirements. To illustrate the application of this rule of thumb, we'll use it to estimate the percentage increase in CPU/command for the comparison shown in this section where we replace the users' A-disks with SFS directories.

Note that the rule of thumb is based on virtual I/Os, not real I/Os. Real I/Os include the effects of minidisk caching and can therefore be much lower. Since not all virtual I/Os are minidisk I/Os and not all minidisks will be migrated to SFS, it would be ideal to have a breakdown of virtual I/Os by minidisk from the current system in order to provide input to the estimation formula. Typically, however, I/Os are only broken down on a DASD volume basis so you need to apply some judgement for volumes where only some of the I/O activity is coming from minidisks that are to be moved to SFS. SEEKS domain monitor data can be used to obtain a breakdown of I/O rate by minidisk for these mixed volumes. You should bear in mind, however, that this is a breakdown of real I/Os, not virtual I/Os.

One good source of virtual I/Os on a volume basis is from the "SSCH+RSCH Plus Avoid Rate" column in VMPRF's DASD_BY_ACTIVITY report. This is the real I/O rate plus the additional I/O rate (if any) that has been avoided by the presence of minidisk caching. In our example, all user A-disks are defined on MDSKnn volumes and there are no additional sources of I/O activity on those volumes. Accordingly, all we need to do is sum the "Plus Avoided" column over all 16 MDSKnn volumes. The result is 218.5 VIOs per second, as shown in Figure 3.

<-----Device------>						<---------SSCH+RSCH-------->			
Num-ber	Volume Serial	Type	Control Unit	Mini-disk Links	On-line Secs	Count	Rate	Plus Avoided	Plus Avoid Rate
FB82	MDSK11	3390-2	3990-3	130	1800	13973	7.8	26515	14.7
F900	MDSK01	3390-2	3990-3	131	1800	13932	7.7	25819	14.3
FB80	MDSK09	3390-2	3990-3	130	1800	13788	7.7	25328	14.1
F905	MDSK06	3390-2	3990-3	129	1800	14197	7.9	25437	14.1
F907	MDSK08	3390-2	3990-3	129	1800	13127	7.3	25008	13.9
FB81	MDSK10	3390-2	3990-3	130	1800	13087	7.3	24903	13.8
F903	MDSK04	3390-2	3990-3	129	1800	13061	7.3	24667	13.7
FB83	MDSK12	3390-2	3990-3	129	1800	13070	7.3	24717	13.7
F904	MDSK05	3390-2	3990-3	129	1800	12779	7.1	24227	13.5
F901	MDSK02	3390-2	3990-3	130	1800	13124	7.3	24358	13.5
F906	MDSK07	3390-2	3990-3	129	1800	12737	7.1	24115	13.4
FB86	MDSK15	3390-2	3990-3	128	1800	13123	7.3	24119	13.4
FB84	MDSK13	3390-2	3990-3	130	1800	13037	7.2	23799	13.2
FB85	MDSK14	3390-2	3990-3	130	1800	12901	7.2	23750	13.2
FB87	MDSK16	3390-2	3990-3	128	1800	12623	7.0	23556	13.1
F902	MDSK03	3390-2	3990-3	129	1800	12390	6.9	23291	12.9
Total VIO/sec moved to SFS -->									218.5

Figure 3. Minidisk volumes from the L29V2101 DASD_BY_ACTIVITY VMPRF report.

Note: *Not all of the report's columns are shown.*

The L29V2101 base measurement ran at 90% processor utilization on a 9121-480. A 9121-480 is roughly a 38 MIPS machine (varies with workload). That means it can execute about 38 million instructions (MI) per wall clock second if the system is running at 100% utilization. In our case, it is running at 90% utilization, so it is executing 0.9*38 = 34 MI per second. 218.5 VIOs/sec / 34 MI/sec = 6.4 VIOs / MI. Applying the rule of thumb, 6% * 6.4 = 38%. This estimate is very close to the 38.3% increase in PBT/CMD(H) that was actually observed.[7]

It is important to note that this method estimates the percentage increase in processor usage per unit work. It does *not* estimate the percentage increase in processor utilization, which also depends upon what happens to the rate at which work is handled by the system.

[7] It is just by chance that the estimated and observed values are in such close agreement in this example. However, this rule of thumb should come within plus or minus 5% of the actual value most of the time. That means, in our example, that the rule of thumb estimates that the true percentage increase should be somewhere between 36% and 40%.

Note that the rule of thumb only applies to minidisk activity that migrated to SFS *filecontrol* directories. Any minidisk activity that is being migrated to SFS dircontrol directories mapped to VM data spaces should be ignored because such a migration normally results in negligible performance differences.

There is a reciprocal relationship between processor usage (per unit work) and the number of users that can be supported from a processor capacity standpoint. In our example, we estimated a 38% increase in processor usage or, equivalently, processor usage will be 1.38 times larger. From that, we can estimate that the system can support 1/1.38 = 0.72 times as many users at the same processor utilization after the proposed migration to SFS — a 28% reduction. This corresponds well with the measured 28.6% reduction.

Migration to VisualGen from CSP

A set of informal performance measurements was obtained to get a general understanding of the performance of migrating an application from CSP to VisualGen.

Two CSP workloads, Heavy Math and List Inquiry, were used for this evaluation. Heavy Math performs a complex mathematical calculation using data entered by the user. This workload does negligible I/O. List Inquiry is an application that retrieves and displays information from a selected VSAM database. Two equivalent VM implementations of these workloads were prepared, one using CSP/AE 3.3 and one using VisualGen Host Services 1.1.0.

Measurements were obtained on a VM/ESA 2.1.0 production system. The CP QUERY TIME command was used to collect processor usage data (TOTCPU), while the CP INDICATE USER command was used to collect virtual I/O count data (IO). These commands were issued before and after each trial run and resource usage was determined by subtraction.

In the setup used to run the CSP applications, actual execution of the application was preceded by two selection menu screens. The processor time required to traverse these screens was measured separately and subtracted from the CSP results. There were no analogous selection menu screens in the VisualGen case.

Table 20. Resource usage comparison between CSP/AE 3.3 and VisualGen 1.1.0

Workload	Metric	CSP	VisualGen	Percent
Heavy Math	CPU (sec)	0.95	0.26	-73
List Inquiry	CPU (sec) Virtual I/Os	1.44 1564	0.65 1163	-55 -26

Note: CPU time is from QUERY TIME (TOTCPU). Virtual I/O is from INDICATE USER (IO).

The Heavy Math results are the average of 5 trials, while the List Inquiry results are the average of 4 trials. Each trial produced similar results.

CSP application code is interpreted while VisualGen application code is compiled. As a result, VisualGen applications tend to use substantially less processor time. Heavy Math shows a larger percentage reduction in processor time than List Inquiry because Heavy Math spends a higher percentage of its time executing application code.

Appendix A. Workloads

The workloads that were used to evaluate VM/ESA 2.2.0 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 21 shows the search-order characteristics of the two environments used for measurements discussed in this document.

Table 21. 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 89 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 22 on page 72.

Table 22. FS8F workload script summary

Script Name	% Used	Script Description
LOGESA	*	Logon and Initialization
WAIT	*	Wait state
CMSSTRT	*	Stagger start of user activity
ASM617F	5	Assemble (HLASM) and Run
ASM627F	5	Assemble and Run
XED117F	5	Edit a VS BASIC Program
XED127F	10	Edit a VS BASIC Program
XED137F	10	Edit a COBOL Program
XED147F	10	Edit a COBOL Program
COB217F	5	COBOL Compile
COB417F	5	Run a COBOL Program
FOR217F	5	VS FORTRAN Compile
FOR417F	5	FORTRAN Run
PRD517F	5	Productivity Aids Session
DCF517F	5	Edit and Script a File
PLI317F	5	PL/I Optimizer Session
PLI717F	5	PL/I Optimizer Session
WND517F	8	Run Windows with IPL CMS
WND517FL	2	Run Windows with LOGON/LOGOFF
HLP517F	5	Use HELP

Note: Scripts with an asterisk (*) in the "% Used" column are run only once each for each user during initialization.

The following are descriptions of each script used in the FS8F workload.

LOGESA: Initialization Script

LOGON userid
SET AUTOREAD ON
IF FS8F0R workload
THEN
  Erase extraneous files from A-disk
  Run PROFILE EXEC to access correct search order,
    SET ACNT OFF, SPOOL PRT CL D, and TERM LINEND OFF
ELSE
  Erase extraneous files from A-directory
  Run PROFILE EXEC to set correct search order, SET ACNT OFF,
    SPOOL PRT CL D, and TERM LINEND OFF
END
Clear the screen
SET REMOTE ON

WAIT: Ten-Second Pause
Leave the user inactive in a 10-second wait loop.

CMSSTRT: Random-Length Pause
Delay, for up to 15 minutes, the start for each user to prevent all users from starting scripts at the same time.

ASM617F: Assemble (HLASM) and Run

QUERY reader and printer
SPOOL PRT CLASS D
XEDIT an assembler file and QQUIT
GLOBAL appropriate MACLIBs
LISTFILE the assembler file
Assemble the file using HLASM (NOLIST option)
Erase the text deck
Repeat all the above except for XEDIT
Reset GLOBAL MACLIBs
Load the text file (NOMAP option)
Generate a module (ALL and NOMAP options)
Run the module
Load the text file (NOMAP option)
Run the module 2 more times
Erase extraneous files from A-disk

ASM627F: Assemble (F-Assembler) and Run

QUERY reader and printer
Clear the screen
SPOOL PRT CLASS D
GLOBAL appropriate MACLIBs
LISTFILE assembler file
XEDIT assembler file and QQUIT
Assemble the file (NOLIST option)
Erase the text deck
Reset GLOBAL MACLIBs
Load the TEXT file (NOMAP option)
Generate a module (ALL and NOMAP options)
Run the module
Load the text file (NOMAP option)
Run the module
Load the text file (NOMAP option)
Run the module
Erase extraneous files from A-disk
QUERY DISK, USERS, and TIME

XED117F: Edit a VS BASIC Program

XEDIT the program
Get into input mode
Enter 29 input lines
Quit without saving file (QQUIT)

XED127F: Edit a VS BASIC Program

Do a FILELIST
XEDIT the program
Issue a GET command
Issue a LOCATE command
Change 6 lines on the screen
Issue a TOP and BOTTOM command
Quit without saving file
Quit FILELIST
Repeat all of the above statements, changing 9 lines instead of 6 and without issuing the TOP and BOTTOM commands

XED137F: Edit a COBOL Program

Do a FILELIST
XEDIT the program
Issue a mixture of 26 XEDIT file manipulation commands
Quit without saving file
Quit FILELIST

XED147F: Edit a COBOL Program

Do a FILELIST
XEDIT the program
Issue a mixture of 3 XEDIT file manipulation commands
Enter 19 XEDIT input lines
Quit without saving file
Quit FILELIST

COB217F: Compile a COBOL Program

Set ready message short
Clear the screen
LINK and ACCESS a disk
QUERY link and disk
LISTFILE the COBOL program
Invoke the COBOL compiler
Erase the compiler output
RELEASE and DETACH the linked disk
Set ready message long
SET MSG OFF
QUERY SET
SET MSG ON
Set ready message short
LINK and ACCESS a disk
LISTFILE the COBOL program
Run the COBOL compiler
Erase the compiler output
RELEASE and DETACH the linked disk
QUERY TERM and RDYMSG
Set ready message long
SET MSG OFF
QUERY set
SET MSG ON
PURGE printer

COB417F: Run a COBOL Program

Define temporary disk space for 2 disks using an EXEC
Clear the screen
QUERY DASD and format both temporary disks
Establish 4 FILEDEFs for input and output files
QUERY FILEDEFs
GLOBAL TXTLIB
Load the program
Set PER Instruction
Start the program
Display registers
End PER
Issue the BEGIN command
QUERY search of minidisks
RELEASE the temporary disks
Define one temporary disk as another
DETACH the temporary disks
Reset the GLOBALs and clear the FILEDEFs

FOR217F: Compile 6 VS FORTRAN Programs

NUCXDROP NAMEFIND using an EXEC
Clear the screen
QUERY and PURGE the reader
Compile a FORTRAN program
Issue INDICATE commands
Compile another FORTRAN program
Issue INDICATE commands
Compile another FORTRAN program
Issue INDICATE command
Clear the screen
Compile a FORTRAN program
Issue INDICATE commands
Compile another FORTRAN program
Issue INDICATE commands
Compile another FORTRAN program
Clear the screen
Issue INDICATE command
Erase extraneous files from A-disk
PURGE the printer

FOR417F: Run 2 FORTRAN Programs

SPOOL PRT CLASS D
Clear the screen
GLOBAL appropriate text libraries
Issue 2 FILEDEFs for output
Load and start a program
Rename output file and PURGE printer
Repeat above 5 statements for two other programs, except
erase the output file for one and do not issue spool printer
List and erase output files
Reset GLOBALs and clear FILEDEFs

PRD517F: Productivity Aids Session

Run an EXEC to set up names file for user
Clear the screen
Issue NAMES command and add operator
Locate a user in names file and quit
Issue the SENDFILE command
Send a file to yourself
Issue the SENDFILE command
Send a file to yourself
Issue the SENDFILE command
Send a file to yourself
Issue RDRLIST command, PEEK and DISCARD a file
Refresh RDRLIST screen, RECEIVE an EXEC on B-disk, and quit
TRANSFER all reader files to punch
PURGE reader and punch
Run a REXX EXEC that generates 175 random numbers
Run a REXX EXEC that reads multiple files of various sizes from both the A-disk and C-disk
Erase EXEC off B-disk
Erase extraneous files from A-disk

DCF517F: Edit and SCRIPT a File

XEDIT a SCRIPT file
Input 25 lines
File the results
Invoke SCRIPT processor to the terminal
Erase SCRIPT file from A-disk

PLI317F: Edit and Compile a PL/I Optimizer Program

Do a GLOBAL TXTLIB
Perform a FILELIST
XEDIT the PL/I program
Run 15 XEDIT subcommands
File the results on A-disk with a new name
Quit FILELIST
Enter 2 FILEDEFs for compile
Compile PL/I program using PLIOPT
Erase the PL/I program
Reset the GLOBALs and clear the FILEDEFs
COPY names file and RENAME it
TELL a group of users one pass of script run
ERASE names file
PURGE the printer

PLI717F: Edit, Compile, and Run a PL/I Optimizer Program

Copy and rename the PL/I program and data file from C-disk
XEDIT data file and QQUIT
XEDIT a PL/I file
Issue RIGHT 20, LEFT 20, and SET VERIFY ON
Change two lines
Change filename and file the result
Compile PL/I program using PLIOPT
Set two FILEDEFs and QUERY the settings
Issue GLOBAL for PL/I transient library
Load the PL/I program (NOMAP option)
Start the program
Type 8 lines of one data file
Erase extraneous files from A-disk
Erase extra files on B-disk
Reset the GLOBALs and clear the FILEDEFs
TELL another USERID one pass of script run
PURGE the printer

WND517F: Use Windows

SET FULLSCREEN ON
TELL yourself a message to create window
QUERY DASD and reader
Forward 1 screen
TELL yourself a message to create window
Drop window message
Scroll to top and clear window
Backward 1 screen
Issue a HELP WINDOW and choose Change Window Size
QUERY WINDOW
Quit HELP WINDOWS
Change size of window message
Forward 1 screen
Display window message
TELL yourself a message to create window
Issue forward and backward border commands in window message
Position window message to another location
Drop window message
Scroll to top and clear window
Display window message
Erase MESSAGE LOGFILE
IPL CMS
SET AUTOREAD ON
SET REMOTE ON

WND517FL: Use Windows with LOGON, LOGOFF

SET FULLSCREEN ON
TELL yourself a message to create window
QUERY DASD and reader
Forward 1 screen
TELL yourself a message to create window
Drop window message
Scroll to top and clear window
Backward 1 screen
Issue a help window and choose Change Window Size
QUERY WINDOW
Quit help windows
Change size of window message
Forward 1 screen
Display window message
TELL yourself a message to create window
Issue forward and backward border commands in window message
Position window message to another location
Drop window message
Scroll to top and clear window
Display window message
Erase MESSAGE LOGFILE
LOGOFF user and wait 60 seconds
LOGON user on original GRAF-ID
SET AUTOREAD ON
SET REMOTE ON

HLP517F: Use HELP and Miscellaneous Commands

Issue HELP command
Choose HELP CMS
Issue HELP HELP
Get full description and forward 1 screen
Quit HELP HELP
Choose CMSQUERY menu
Choose QUERY menu
Choose AUTOSAVE command
Go forward and backward 1 screen
Quit all the layers of HELP
RELEASE Z-disk
Compare file on A-disk to C-disk 4 times
Send a file to yourself
Change reader copies to two
Issue RDRLIST command
RECEIVE file on B-disk and quit RDRLIST
Erase extra files on B-disk
Erase extraneous files from A-disk

IBM Office Benchmark (IOB)

Workload Description

The IBM Office Benchmark (IOB) Version 2.1 is a corporate-wide benchmark designed to measure generic office system performance. It consists of the office user definition, the databases for calendars, the documents, the mail, and the work the office users do.

The IOB measurements included in this report use the DisplayWrite*/370 2.1.0 and the OfficeVision/VM licensed programs.

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 and reach the OfficeVision main menu (the A00 screen).
- After a random delay of up to 10 minutes, each user selects a script and starts. (The random delay prevents all users from starting at once).
- A stabilization period (45 minutes) 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 (30 minutes).
- At the end of the measurement interval, the performance data are reduced and analyzed.

The IOB workload does not aim for a specific think time or use a certain think time distribution. Instead, the think time is dictated by the IOB workload. The think time includes an average two second delay between commands issued by TPNS, the built-in think times that are part of the IOB scripts, and the IOB script scheduling algorithm. When users finish running a script, the script scheduling algorithm calculates how much time was spent running the script, subtracts this number from ten minutes, and delays the user for the resulting amount of time. Thus, if a script completed in 7.9 minutes, the user would be delayed for 2.1 minutes before starting the next script and this time would be included in the user's think time.

IOB Script Descriptions

The IOB workload consists of nine scripts (scenarios). These scripts are listed in Table 23 with their defined use factor.

Table 23. IOB workload script summary

Script Name	% Used	Script Description
VMB2LML	17	Send Note and Process Light Mail
VMB2HML	17	Send Note and Process Heavy Mail
VMB2VCAL	13	View Individual Calendar
VMB2UCAL	13	Update Individual Calendar
VMB2DIR	20	View User Directory
VMB2CDOC	7	Create Small Text Document
VMB2UDOC	7	Revise Small Text Document
VMB2EB	3	Enter/Exit Office
VMB2ONOF	3	Logoff/Logon System

The following is the list of tasks in each script within the IOB workload.

Send Note and Process Light Mail

- Create a note and send the note to two users.
- View the note log.
- View the first item, a note.
- Delete the first item, a note.
- Open Mail and View the In-Basket.
- View the first item, a note.
- Delete the first item, a note.

Send Note and Process Heavy Mail

- Create a note and send the note to two users.
- View the note log.
- View the first item, a note.
- Delete the first item, a note.
- Open Mail and View the In-Basket.
- View the first item, a note.
- Forward the first item to another user with an attachment.
- Delete the original first item, a note.
- View the eighth item in the mail list, a two page document.
- Print the document.

View Individual Calendar

- View the user's calendar for Wednesday of a defined week.

Update Individual Calendar

- View the user's calendar for Wednesday of a defined week.
- Delete a meeting.
- Add a meeting.

View User Directory

- Search the user directory based on a random user name and view the person's telephone number.

Create Small Text Document

- Get a pre-stored document format.
- Key in a two-page document.
- Save the document.
- Print the document.
- Delete the document.

Revise Small Text Document

- Open a two-page document for revision.
- Move one paragraph.
- Delete one paragraph.
- Insert one paragraph.
- Save the altered document.
- Send the document to three users.

Enter/Exit Office

- Enter the office software environment.
- Exit the office software environment.

Logon/Logoff System

- Log off from the system.
- Log back onto the system and enter the office environment.

VSE Guest (DYNAPACE)

Workload Description

PACE is a synthetic VSE batch workload consisting of 7 unique jobs representing the commercial environment. This set of jobs is replicated 16 times, producing the *DYNAPACE* workload. The first 8 copies run in 8 static partitions and another 8 copies run in 4 dynamic classes, each configured with a maximum of 2 partitions. The 7 jobs are:

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.

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 8 copies of PACE. Four dynamic classes, each with 2 partition assignments, run another 8 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
- 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.

Appendix B. Configuration Details

Saved Segments

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 CMS and OV/VM environments in this report used the following saved segments:

CMS	Contains the CMS nucleus and file status tables (FSTs) for the S- and Y-disks.
CMSFILES	Contains the SFS server code in the DMSDAC and DMSSAC logical segments.
CMSPIPES	Contain CMSPIPES code in the PIPES logical segment.
CMSINST	Contains the execs-in-storage segment.
CMSVMLIB	Contains the following logical segments: • VMLIB contains the CSL code. • 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.
OFSSEG	Contains OV/VM user functions
EPUYSSEG	Contains OV/VM mailbox manager code
DW370210	Contains the DW370 module
DDDCL210	Contains the DW370 compiled CLISTS
DW362	Contains FSTs for the DW/370 362 disk
ADM399	Contains FSTs for the OV/VM 399 disk
GCSXA	Contains the GCS nucleus.
VTAMXA	Contains the VTAM code.

Server Options: SFS DMSPARMS

This section lists the start-up parameter settings used by each of the SFS servers. The start-up parameters determine the operational characteristics of the file pool server. The SFS servers used the following DMSPARMS file:

```
ADMIN MAINT U3 OPERATOR MARK
NOBACKUP
FULLDUMP
FILEPOOLID fp_name
NOFORMAT
ACCOUNT
MSGS
SAVESEGID CMSFILES
USERS  nnnn
```

For all SFS measurements, the SAVESEGID is specified to identify the segment containing the file pool server runnable code. USERS was set equal to the number of logged on users that were connected to the SFS file pool server during the measurement. The USERS parameter is used by the SFS server to configure itself with the appropriate number of user agents and buffers.

Server Options: 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 the SFS and CRR tuning parameters, see the *CMS File Pool 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 24 covers all performance measurement results except for migration results, which are covered by Table 11 on page 41. 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*

220 *VM/ESA Version 2 Release 2.0 Performance Report* (this document)

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

Table 24 (Page 1 of 3). Sources of VM performance measurement results

Subject	Report(s)
Migration	see page 41
New Functions	
Coordinated Resource Recovery	10
VM Data Spaces (Use by SFS)	11
3990-3 DASD Fast Write Support	11
CMS Pipelines	11
Inter-System Facility for Communications (ISFC)	11 22
ECKD* Support	11
FBA DASD Support	20
CP Configurability	20
DIAGNOSE Code X’250’	20
Extended CMS File System Interfaces	20
Virtual Disk in Storage	21 22
Load Wait State PSW Improvements	21
REXX SAA* Level 2 Architecture	21
Minidisk Cache Enhancements	22
Share Capping and Proportional Distribution	22
SPXTAPE Command	22
ISFC Changes	22
POSIX	210
DCE	210
GCS TSLICE Option	210
Removal of VMCF Master Processor Dependency	220

Table 24 (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
TCP/IP 2.4.0	220
OfficeVision 1.3.0	220
OfficeVision SFS A-Directory Support	220
Migration to VisualGen from CSP	220
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
Configuration of Storage as Central or Expanded	220
Processor Capacity	
3090-600J	10
9121-480	11
9021-720	11
9021-900	20
9121-742	21
9021-941	22
PC Server 500	210
PC Server 520	220
RS/6000 591	220
9672-R53	220

Table 24 (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.

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:

$$10 \times \frac{(\text{TOTAL} - \text{TOTAL EMUL})}{\text{ETR (T)}}$$

For the PACE workload:

$$\text{PBT/CMD} - \text{EMUL/CMD}$$

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:

$$\text{PBT/CMD (H)} - \text{EMUL/CMD (H)}$$

CP CPU/CMD (V) Server. CP processor time, in milliseconds, run in the designated server machine per command. This is calculated by:

$$\left(\frac{1}{V\_Time \times \text{ETR (T)}}\right) \times \sum_{s \in \text{ server class}} (TCPU_s - VCPU_s)$$

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 X04 instructions used per command. DIAGNOSE code X04 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 X08 instructions used per command. DIAGNOSE code X08 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 X0C instructions used per command. DIAGNOSE code X0C 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*_0*C* 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 X10 instructions used per command. DIAGNOSE code X10 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 X14 instructions used per command. DIAGNOSE code X14 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 X58 instructions used per command. DIAGNOSE code X58 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 X98 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 X98 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 XA4 instructions used per command. DIAGNOSE code XA4 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 XA8 instructions used per command. DIAGNOSE code XA8 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 X214 instructions used per command. DIAGNOSE code X214 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 X268 instructions used per command. DIAGNOSE code X268 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 270/CMD. The number of DIAGNOSE code X270 instructions used per command. DIAGNOSE code X270 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. Its output is the same as DIAGNOSE code X0C with two additional fields that provide the date as mm/dd/yyyy and yyyy-mm-dd. This diagnose interface was added in VM/ESA 2.2.0 as part of the year 2000 support. This is calculated by:

$$\frac{DIAG\_270}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG_*270 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.

FCON/ESA. FCON/ESA is a program that is available from IBM that provides performance monitoring capabilities with system console operation in full screen mode. FCON/ESA can provide an immediate view of system performance or post process its own history files or VM/ESA monitor data for selected data. Threshold monitoring and user loop detection is provided. FCON/ESA also has the ability to monitor remote systems.

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. The CMS control block that contains information about a file belonging to a minidisk or SFS directory.

GB. Gigabytes. 1024 megabytes.

GUEST SETTING. This field represents the type of VSE guest virtual machine in a PACE measurement. This fields possible values are V=V, V=F or V=R.

GUESTWT/CMD. The number of entries into guest enabled wait state per job. This is calculated by:

$$60 \times \frac{\text{GUESTWT/SEC}}{\text{ETR (H)}}$$

GUESTWT/SEC. The number of entries into guest enabled wait state per second.

This field is taken from the NSEC column for the RTM total count since last reset, for the GUESTWT field in the RTM SYSTEM screen.

Hardware Instrumentation. See Processor Instrumentation

HT5. One of the CMS-intensive workloads used in the Large Systems Performance Reference (LSPR) to evaluate relative processor performance.

IML MODE. This is the hardware IML mode used in VSE guest measurements. The possible values for this field are 370, ESA, or LPAR.

Instruction Path Length. The number of machine instructions used to run a given command, function or piece of code.

Internal Response Time. The response time as seen by CP. This does not include line or terminal delays.

IO TIME/CMD (Q). Total elapsed time in seconds spent doing SFS file I/Os per command. This is calculated by:

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

Total_BIO_Request_Time is from the QUERY FILEPOOL STATUS command.

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

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

Total_IO_Requests is from the QUERY FILEPOOL STATUS command.

ISFC. Inter-System Facility for Communications

ITR. Internal Throughput Rate. This is the number of units of work accomplished per unit of processor busy time in an nonconstrained environment. For the FS7F, FS8F, and VSECICS workloads this is represented as commands per processor second. For the PACE workload, this is represented as jobs per processor minute. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads, this is found from the TOTALITR for SYS_ITR on the RTM TRANSACT screen.

For the PACE workload:

$$100 \times \frac{\text{ETR (H)}}{\text{UTIL/PROC}}$$

ITR (H). See ITR. This is the hardware based measure. In this case, ITR is measured in external commands per unit of processor busy time. For the FS7F, FS8F, and VSECICS workloads this is represented as commands per processor second, while for the PACE workload this is represented in jobs per processor minute. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

$$100 \times \frac{\text{ETR (T)}}{\text{UTIL/PROC (H)}}$$

For the PACE workloads:

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

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.

ITR (V). See ITR. This is the VMPRF-based measure. ITR is measured in external commands per unit of processor busy time. This is calculated by:

$$100 \times \frac{\text{ETR (T)}}{\text{UTIL/PROC (V)}}$$

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.

ITRR (V). See ITRR. This is the ITR (V) normalized to a specific run. This is calculated by:

$$\frac{\text{ITR (V)}}{\text{ITR (V)}_1}$$

ITR (V)_1 is the ITR (V) of the first run in a given table.

IUCV. Inter-User Communication Vehicle. 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:

$$\sum_{f \in \text{ filepools}} \textit{LUW\_Rollbacks}_f$$

LUW_Rollbacks is from the QUERY FILEPOOL STATUS command.

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

This is taken from the %EM column for the first processor listed in the LOGICAL CPU STATISTICS section of the RTM CPU screen. The total RTM interval time value is used (<-..).

MASTER EMUL (H). Total emulation state utilization for the master processor. For uniprocessors this is the same as TOTAL EMUL and is generally not shown. This is the hardware based calculation.

This is taken from the %CPU column of the GUES-CPn line of the REPORT file for the master processor number as shown by RTM. In RTM, the first processor listed on the CPU screen is the master processor.

MASTER TOTAL. Total utilization of the master processor. For uniprocessor this is the same as TOTAL and is generally not shown.

This is taken from the %CPU column for the first processor listed in the LOGICAL CPU STATISTICS section of the RTM CPU screen. The total RTM interval time value is used (<-..).

MASTER TOTAL (H). Total utilization of the master processor. For uniprocessor this is the same as TOTAL (H) and is generally not shown. This is the hardware based calculation.

This is taken from the %CPU column of the SYST-CPn line of the REPORT file for the master processor number as shown by RTM. In RTM, the first processor listed on the CPU screen is the master processor.

MB. Megabytes. One megabyte is 1,048,576 bytes.

MDC AVOID. The number of DASD read I/Os per second that were avoided through the use of minidisk caching.

For VM releases prior to VM/ESA 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.

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:

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

For the PACE workload:

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

TOTALCNT is the count for the individual privop taken over the total RTM time interval on the RTM PRIVOPS Screen. *RTM_Time* is the total RTM time interval taken from the RTM PRIVOPS screen. **Note**: PRIVOPS are recorded differently in 370 and XA modes.

PRIVOPS (Privileged Operations). See Privileged Operation.

Processor Instrumentation. An IBM* internal tool used to obtain hardware-related data such as processor utilizations.

Processor Utilization. The percent of time that a processor is not idle.

Processors. The data field denoting the number of processors that were active during a measurement.

This is from the NC field under CPU statistics on the RTM CPU screen.

PSU. Product Service Upgrade

Production File Pool. An SFS file pool in which users are enrolled with space. All SFS read/write activity is to production file pools.

QUICKDSP ON. When a virtual machine is assigned this option, it bypasses the normal scheduler algorithm and is placed on the dispatch list immediately when it has work to do. It does not spend time in the eligible lists. QUICKDSP can be specified either via a CP command or in the CP directory entry.

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.

Run ID. An internal use only name used to identify a performance measurement.

SAC Calls / FP Request. The average number of calls within the SFS server to its Storage Access Component (SAC) per file pool request. In environments where there are multiple file pools, this average is taken over all file pool servers. This is calculated by:

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

Sac_Calls and *Total_Filepool_Requests* are from the QUERY FILEPOOL STATUS command.

Seconds Between Checkpoints. The average number of seconds between SFS file pool checkpoints in the average file pool. This is calculated by:

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

Checkpoints_Taken is from the QUERY FILEPOOL STATUS command.

SET RESERVED (Option). This is a CP command that can be used to allow a V=V virtual machine to have a specified minimum number of pages resident in real storage. It is used to reduce paging and improve performance for a given virtual machine.

SFSTIME. The elapsed time in seconds between QUERY FILEPOOL STATUS invocations for a given file pool done at the beginning and end of a measurement.

SFS TIME/CMD (Q). Total elapsed time per command, in seconds, required to process SFS server requests. This is calculated by:

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

Filepool_Request_Service_Time is from the QUERY FILEPOOL STATUS command.

SHARE. The virtual machine's SHARE setting. The SET SHARE command and the SHARE directory statement allow control of the percentage of system resources a virtual machine receives. These resources include processors, real storage and paging I/O capability. A virtual machine receives its proportion of these resources according to its SHARE setting. See Relative and Absolute Share.

Shared Storage. The portion of a virtual machines storage that is shared among other virtual machines (such as saved segments). This is usually represented in pages.

SHRPGS. The number of shared frames currently resident.

SIE. ESA Architecture instruction to Start Interpretive Execution. This instruction is used to run a virtual machine in emulation mode.

SIE INTCPT/CMD. The number of exits from SIE which are SIE interceptions per command or job. SIE is exited either by interception or interruption. An intercept is caused by any condition that requires CP interaction such as I/O or an instruction that has to be simulated by CP. This is calculated by:

$$\frac{Percent\_Intercept \times \text{SIE/CMD}}{100}$$

Percent_Intercept is taken from the %SC field for average of all processors for the total RTM time interval (<-..) on the RTM CPU screen.

SIE/CMD. SIE instructions used per command or job. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

$$\frac{SIE\_SEC}{\text{ETR (T)}}$$

For the PACE workload:

$$60 \times \frac{SIE\_SEC}{\text{ETR (H)}}$$

SIE_SEC is taken from the XSI field for the total for all processors for the total RTM time interval (<-..) on the RTM CPU screen.

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.

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:

TOT INT × ETR RATIO

TOT PAGES/USER. The total number of pages that are associated, on average, with each end user virtual machine. This is taken from VMPRF report UCLASS_RESOURCE UTIL and is the sum of resident storage pages, expanded storage pages, and DASD page slots for the ″Users″ class. This is a measure of how many unique pages are touched during execution of the workload by the average end user.

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:

UTIL/PROC (H) × 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:

$$\frac{\text{TOTAL (H)}}{\text{TOTAL EMUL (H)}}$$

T/V Ratio. See TVR

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.

UTIL/PROC (V). Average utilization per processor reported VMPRF.

This is taken from the CPU Pct Busy field in the VMPRF SYSTEM_SUMMARY_BY_TIME report.

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:

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

For the PACE workload:

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