Virtual Machine/
Enterprise Systems Architecture
Performance Report
Version 2 Release 3

March 19, 1998

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.

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

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM licensed program in this publication is not intended to state or imply that only IBM′s program may be used. Any functionally equivalent 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 2.3.0 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 2.3.0. See the IBM Programming Announcement for VM/ESA 2.3.0 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
DisplayWrite
ES/9000
IBM
MQ
MQSeries
MVS
OfficeVision
OpenEdition
OS/2
PR/SM
Processor Resource/Systems Manager
RAMAC
RS/6000
System/390
S/390
Virtual Machine/Enterprise Systems Architecture
VM/ESA
VM/XA
VSE/ESA
VTAM
3090

Acknowledgements

The following people contributed to this report:

Bill Bitner
Wes Ernsberger
Bill Guzior

Abstract

The *VM/ESA Version 2 Release 3.0 Performance Report* summarizes the performance evaluation of VM/ESA 2.3.0, including TCP/IP Function Level 310. Measurements were obtained for CMS-intensive, VSE guest, Telnet, and FTP environments on various ES/9000 processors. Discussion covers the performance changes in VM/ESA 2.3.0, the performance effects of migrating from VM/ESA 2.2.0 to VM/ESA 2.3.0, the performance effects of migrating from TCP/IP 2.4 to TCP/IP 310, the performance of new functions provided in VM/ESA 2.3.0, and additional evaluations.

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
- *C for VM/ESA Library Reference,* SC23-3908
- *CMS Application Development Guide,* SC24-5450
- *MQSeries: Application Programming Guide* SC33-0807

The following publications are performance reports for earlier VM/ESA releases. The topics covered in these reports and this report are listed in the VM/ESA Performance Report Directory at http://www.vm.ibm.com/perf/perfdir.html.

- *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
- *VM/ESA Version 2 Release 2.0 Performance Report,* http://www.vm.ibm.com/perf/docs/

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 3.0, including TCP/IP Function Level 310. Measurements were obtained for the CMS-intensive, VSE guest, Telnet, and FTP environments on various 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 2.3.0 includes a number of performance enhancements (page 3). Some changes have the potential to adversely affect performance (page 6). Lastly, a number of changes were made that affect VM/ESA performance management (page 7).

Migration from VM/ESA 2.2.0 and TCP/IP 2.4

Benchmark measurements show the following performance results for VM/ESA 2.3.0 relative to VM/ESA 2.2.0:

CMS-intensive Overall performance was similar to VM/ESA 2.2.0. The internal throughput rate (ITR) decreased by an average 0.5% while response time improved by an average of 1% (page 13).

VSE guest ITR decreased by an average 0.9% for the measured V=R and V=V environments, while elapsed times were equivalent (page 29).

For the measured Telnet and FTP environments, TCP/IP 310 processor usage in the TCPIP virtual machine decreased by 3.7% relative to TCP/IP 2.4. This improvement resulted from the pathlength improvements made to the TCP/IP stack in this release (page 33).

Migration from Other VM Releases: The performance measurement data in this report can be used in conjunction with similar data in the previous VM/ESA performance reports to get a general understanding of the performance aspects of migrating from earlier VM releases to VM/ESA 2.3.0 (page 39).

New Functions

Minidisk caching now has a record level option for use in unusual cases where the default full track minidisk caching is not appropriate. Results illustrate that record level caching is most suitable for applications that randomly access small amounts of data in very large files (page 44).

Fork support: Processor usage measurements are provided for the fork(), execl(), and wait() OpenEdition functions (page 48).

MQSeries* client support: On a 2003-156, application virtual machine processor usage per MQGET or MQPUT was measured to be about 1 millisecond for a C application and about 14 milliseconds for a REXX application (page 49).

Additional Evaluations

Measurement results show similar response times and increased processor usage (8% to 12%) when migrating a CMS-intensive workload from VTAM to TCP/IP Telnet (page 51).

Measurement results demonstrate that Telnet performance is no longer adversely affected by the presence of unused TCP/IP buffers. This results from improvements that were made in this area (page 58).

The maximum throughput that can be handled by any given SMTP server virtual machine has been substantially increased with TCP/IP 310. For the measured workload and system configuration, throughput increase multiples ranging from 1.6 to 3.4 relative to TCP/IP 2.4 were observed (page 62).

Overall performance of OfficeVision 1.4.0 decreased slightly relative to OfficeVision 1.3.0. External response time increased by 5%, while internal throughput decreased by 1.1%. These differences should be appreciably smaller when running with the compiled version of the REXX execs supplied with OfficeVision (page 67).

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 71).

Changes That Affect Performance

This chapter contains descriptions of various changes to VM/ESA 2.3.0 that affect performance. This information is also available at http://www.vm.ibm.com/perf, along with corresponding information for previous releases.

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 6. 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. Those results serve to further illustrate where these changes apply and how they may affect performance.

"Performance Management" on page 7 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
  - Reduced Segment Table Storage
  - Record Level Minidisk Cache
  - Improved Pacing for Secondary User Console Output
- CMS
  - Execute Macro Rewrite
  - CMSINST Shared Segment Additions
  - PEEK Improvements
- TCP/IP
  - Reduced TCP/IP Processor Usage
  - Reduced TCP/IP Real Storage Usage
  - TCP/IP RFC 1323
  - Increased SMTP Capacity

Reduced Segment Table Storage

For each virtual address space, CP must maintain a contiguous, page-aligned segment table in fixed storage for the hardware to use to translate virtual addresses into real addresses. In prior releases, the storage allocated for any given primary address space segment table was always one of three possible sizes. Segment tables of 32M or less were allocated from space reserved at the beginning of that virtual machine's VMDBK and therefore took no additional storage. Primary address spaces greater than 32MB but not exceeding 1GB had their segment table allocated from the beginning of a separate 4K page. Finally, primary address spaces greater than 1GB and up to the architechted maximum of 2GB had their segment table allocated from the beginning of two contiguous

4K pages. In either of these last two cases, any remaining space not needed for the segment table was unused.

In VM/ESA 2.3.0, this space is now available for satisfying other user free storage requests. As a result, fixed storage requirements go up in a much more continuous manner as segment table size increases. Although it is still best for performance if the segment tables are small, it is no longer important to try to keep the virtual machine segment table size below 32MB or, failing that, below 1GB. The net effect of this change is to simplify shared segment and virtual machine size management, and to improve performance in cases where large numbers of segments or virtual machines exceed 32MB or 1GB.

Installations that currently define their shared segments top-down starting at 1GB should consider relocating them to lower virtual address ranges in order to benefit from this change.

Record Level Minidisk Cache

This support, first available on VM/ESA 2.1.0 and VM/ESA 2.2.0 as APAR VM61045, has now been integrated into VM/ESA 2.3.0. It is intended for unusual cases where the default full track minidisk caching is not appropriate. This typically occurs when a large database (hundreds of cylinders) is implemented as a single CMS file and the application does large numbers of random accesses to small amounts of data in that file. With full track minidisk caching, this can result in serious cache thrashing in the DASD control unit, expanded storage, and main storage. In such cases, using record level minidisk cache for that file′s minidisk is likely to improve performance.

This support is limited to:

- 4KB-formatted CMS minidisks on non-FBA DASD
- I/O done by diagnose X′18′, diagnose X′A4′, diagnose X′250′, or the *BLOCKIO CP system service

Record level caching can be enabled by use of the RECORDMDC option with the SET MDCACHE MDISK command or the MINIOPT directory control statement. The current setting can be determined by using the QUERY MDCACHE MDISK command.

See “Record Level Minidisk Caching” on page 44 for measurement results. See the *VM/ESA: Performance* book for minidisk cache tuning guidelines.

Improved Pacing for Secondary User Console Output

The pacing algorithm for console output directed to a secondary user (established using the CP SET SECUSER command) has been improved. In prior releases, the pacing algorithm limited output to 22 lines per second. Once this limit was reached, the virtual machine generating the output would be suspended for a second. If this virtual machine is a server doing tracing or otherwise generating a high rate of console I/O, that server′s performance could be degraded. In VM/ESA 2.3.0, this problem has been reduced because the limit has been raised to 255 lines per second.

Execute Macro Rewrite

The EXECUTE Xedit macro was rewritten for improved performance and maintainability. This resulted in a 6% CPU usage reduction for EXECUTE when used by FILELIST and a 13% CPU usage reduction for EXECUTE when used by RDRLIST. EXECUTE macro CPU usage is essentially unchanged when used by DIRLIST, CSLLIST, or MACLIST.

CMSINST Shared Segment Additions

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

```
ALL       XEDIT  (ALL Xedit command)
APILOAD   EXEC   (add REXX copy files for multitasking)
DEFAULTS  EXEC   (setup for productivity aids)
EXECUPDT  EXEC   (apply updates to an executable)
HELP      XEDIT  (HELP in Xedit)
JOIN      XEDIT  (JOIN in Xedit)
MOREHELP  EXEC   (HELP for more information)
OPENVM    EXEC   (front end to OPENVM)
PREFIXX   XEDIT  (X prefix command)
PRFSHIFT  XEDIT  (> prefix command)
PRFSHOW   XEDIT  (S prefix command)
PROFTMPL  XEDIT  (profile for templates)
RECEIVE   XEDIT  (RECEIVE from PEEK)
RGTLEFT   XEDIT  (PF10 in default Xedit)
SPLTJOIN  XEDIT  (PF11 in default Xedit)
X$EUPD$X  XEDIT  (EXECUPDT with NOCOMMENT option)
X$EXCM$X  EXEC   (EXECMAP)
X$LKED$X  XEDIT  (LKED)
X$TMPL$X  XEDIT  (TEMPLATE option with CSLLIST)
```

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.

PEEK Improvements

A number of changes were made to the PEEK command for improved function and performance. The performance results depend upon the format of the file being viewed. Measurements showed a 17% CPU usage reduction for print files, a 15% reduction for punch files, and a 25% reduction for disk dump files. There was no significant change in the performance of PEEK for NETDATA files.

Reduced TCP/IP Processor Usage

Processor usage by the TCPIP virtual machine has been reduced by 3.7% for the measured environments. An additional processor usage improvement (up to 1%) can be realized on processors that support the checksum (CKSM) instruction.

See “TCP/IP” on page 33 for measurement results.

Reduced TCP/IP Real Storage Usage

In prior TCP/IP releases, an increase in the TCP/IP buffer pool sizes resulted in an increase in TCPIP virtual machine real storage usage, even when the additional buffers were never used. With TCP/IP 310, this has been corrected so that excess buffers have no appreciable effect on TCPIP real storage requirements. This means that you can now provide extra buffers without the risk of undesirable performance effects.

See "Reduced TCP/IP Real Storage Requirements" on page 58 for measurement results.

TCP/IP RFC 1323

TCP/IP 310 implements RFC 1323, a TCP protocol extension that allows window sizes exceeding 64KB to be negotiated. This RFC applies to high bandwidth, high latency connections. In such cases, it can increase maximum throughput by increasing the amount of data that can be transmitted before an acknowledgement from the receiving system is required.

Increased SMTP Capacity

The SMTP server has been redesigned so that it does fewer minidisk I/Os and reads spool files asynchronously (using the *SPL CP system service). These changes have resulted in substantial improvements to the maximum throughput that one SMTP server virtual machine can deliver. Example measurements show 1.3-fold to 3.4-fold throughput increases relative to TCP/IP 2.4. The actual degree of improvement that will be experienced on a given system configuration will primarily depend upon where the SMTP log file is written (spool or minidisk), average access time to the DASD volume containing the SMTP A-disk, the average access time to the spool volumes, processor availability, and processor speed. The largest relative improvements will be observed on systems where the log file is written to spool and that are characterized by long DASD access times, high processor availability, and high processor speed.

See "Improved SMTP Server Capacity" on page 62 for measurement results.

Performance Considerations

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

- DASD I/O Queue Ordering Algorithm Removed
- Potential Shared Segment Overlaps
- NFS Performance

DASD I/O Queue Ordering Algorithm Removed

In prior releases, CP ordered DASD I/O requests in an attempt to minimize seek time. This has been made obsolete and, in some cases, even counterproductive by the newer DASD technologies. In addition, in certain unusual cases, this algorithm can result in very long service delays for a given user. Because of these considerations, this algorithm has been disabled.

For the great majority of cases, this change is expected to have either no discernible effects or result in improved system characteristics. However, DASD I/O access times might increase in some situations. This is more likely on systems with substantial DASD I/O queueing, old DASD, and non-cached control units.

Potential Shared Segment Overlaps

The portion of the CMS saved system that resides above the 16MB line has been extended by one megabyte and now ends at location X′13FFFFF′. While installing VM/ESA 2.3.0, check to make sure that this has not caused any overlaps with other shared segments in your system. In addition, if you have any non-relocatable modules that were generated to load between X′1300000′ and X′13FFFFF′, they will need to be regenerated to run in another location.

NFS Performance

The performance of NFS when used with shared file system or byte file system directories is significantly less than the performance of NFS when used with CMS minidisks.

Performance Management

These changes affect the performance management of VM/ESA and TCP/IP VM.

- Monitor Enhancements
- Improved MONVIEW Performance
- NETSTAT 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 2.3.0.

- System Configuration Data

  Various fields were added to the system configuration data record (domain 1 record 4). These include the volume serial numbers for the checkpoint and warm start areas and the system identifier name.

- IUCV Connection Information

  Two new fields were added to the user activity data record (domain 4 record 3). These fields are the maximum number of IUCV connections allowed and the current number of IUCV connections in use. The maximum number of connections is from the system default or the MAXCONN setting in the user directory entry.

- Improved MDC Counter

  The MDC read request count in the expanded storage data record (domain 0 record 14) is now more accurate. In previous releases, it was possible for this number to be inaccurate in either direction. This field is often used to compute MDC hit ratios.

- APPLDATA (domain 10) Interface

  Enhancements were made to diagnose X′DC′, which is the interface applications can use to contribute data to the monitor. Applications can now easily contribute event and configuration data to the APPLDATA domain.

APPLDATA monitor records are now being contributed by two TCP/IP servers—TCPIP and TFTPD—to allow for improved monitoring of their performance.

- TCPIP

  The TCP/IP 310 stack machine can now generate APPLDATA event and sample records. The layouts of these records are provided in the *Performance* manual. The following types of records are produced:

```
TYPE    Rec  Description

Sample  00   TCP/IP MIB Record
Event   01   TCP/IP TCB Open Record
Event   02   TCP/IP TCB Close Record
Config  03   TCP/IP Pool Limit Record
Sample  04   TCP/IP Pool Size Record
Sample  05   TCP/IP LCB Record
Event   06   TCP/IP UCB Open Record
Event   07   TCP/IP UCB Close Record
Config  08   TCP/IP Link Definition Record
Sample  09   TCP/IP ACB Record
Sample  0A   TCP/IP CPU Record
Event   0B   TCP/IP CCB Record
Sample  0C   TCP/IP Tree Size Record
Config  0D   TCP/IP Home Record
```

These records are provided if the APPLMON option is specified for the TCPIP virtual machine, subject to the new MONITORRECORDS statement in the PROFILE TCPIP configuration file:

```
NORECORDS      no monitor records (default)
MOSTRECORDS    all records except for the ACB records
ALLRECORDS     all monitor records
```

- TFTPD

  The TFTPD server was first introduced as an APAR to TCP/IP 2.4 as part of VM/ESA's support for the IBM Network Station. TFTPD contributes APPLDATA sample records to CP monitor. These records are always provided if the APPLMON option is specified for the TFTPD virtual machine. They can be used to get a more detailed understanding of the work that is being performed by the TFTPD server. The information provided includes:

  files read by client
  bytes read
  files read from cache
  elapsed time spent reading files
  files written to the TFTPD server
  bytes written
  elapsed time spent writing files
  information on failed transactions
  timeouts waiting for an acknowledgement

Improved MONVIEW Performance

The MDATPEEK stage, which is part of the MONVIEW tool that is shipped on the 3B2 samples disk was rewritten from REXX to assembler. This greatly improves MONVIEW performance for certain cases. MONVIEW can be used to view CP monitor data. For more information on MONVIEW, see the MONVIEW SAMPLIST file on the samples disk.

NETSTAT Enhancements

NETSTAT has been improved to have greater selectability, limiting the amount of output to what you actually need. In addition, the NETSTAT INTERVAL command now has a fullscreen interface where the data can be scrolled and sorted by field.

Effects on Accounting Data

None of the VM/ESA 2.3.0 performance changes are expected to have a significant effect on the values reported in the virtual machine resource usage accounting record.

VM Performance Products

This section contains information on the support for VM/ESA 2.3.0 provided by VMPRF, RTM/ESA, FCON/ESA, and VMPAF.

VM Performance Reporting Facility 1.2.1 (VMPRF) will run on VM/ESA 2.3.0 with the same support as VM/ESA 2.2.0. The latest service is recommended.

Realtime Monitor VM/ESA 1.5.2 (RTM/ESA) requires APAR GC05430 (PTF UG03868) to run on VM/ESA 2.3.0. RTM/ESA has been updated to use a field new to VM/ESA 2.3.0 when calculating the minidisk hit ratio. With this new field, the calculation corresponds more closely to the value returned from the CP INDICATE command. RTM/ESA will continue to do the old calculation when running on earlier VM/ESA releases.

FCON/ESA Versions 2.3.02 and 3.1.00 will run on VM/ESA 2.3.0 with the same support as VM/ESA 2.2.0.

Performance Analysis Facility/VM 1.1.3 (VMPAF) will run on VM/ESA 2.3.0 with the same support as VM/ESA 2.2.0.

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
- 9121-480
- 9121-320

  To run as a 9121-320, one processor was varied offline from the 9121-480 hardware configuration screen.

Software

A pre-GA (General Availability) level of VM/ESA 2.3.0 was used for the measurements in this report.

Other VM/ESA releases were measured for this report. VM/ESA 2.2.0 was at the GA+first-RSU (Recommended Service Upgrade) level. The service that was part of VM/ESA 2.2.0 after the first RSU level and integrated into VM/ESA 2.3.0 can account for some of the performance differences between VM/ESA 2.2.0 and VM/ESA 2.3.0.

See the appropriate workload section in Appendix A, "Workloads" on page 78 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 78.
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 88 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 12.

   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 2.2.0 and TCP/IP 2.4

This chapter examines the performance effects of migrating from VM/ESA 2.2.0 to VM/ESA 2.3.0 and from TCP/IP 2.4 to TCP/IP Function Level 310. The following environments were measured: CMS-intensive, VSE guest, Telnet, and FTP.

CMS-Intensive

The VM/ESA 2.3.0 performance results are similar to the corresponding VM/ESA 2.2.0 results for the three measured CMS-intensive environments. ITR decreases ranging from 0.3% to 0.7% and external response time changes ranging from a 3% decrease to a 1% increase were observed.

Although VM/ESA 2.3.0 includes a number of performance improvements (see “Performance Improvements” on page 3), none of them apply in a significant way to the measured CMS environments.

Measurement results and discussion for each of these three 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)
Tape: 3480 (Monitor)

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*- Number of Volumes -*					
			PAGE	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-3	RAMAC 2	4		13		32		
3390-2	3990-3	4	6		7			2 R
3390-2	3990-2	4	16		6			

Note: *R* next to the DASD counts means basic cache enabled; *W* means DASD fast write (and basic cache) enabled. RAMAC 2 refers to the RAMAC 2 Array Subsystem with 256MB cache and drawers in 3390-3 format. With the RAMAC 2 Array Subsystem, read and write caching are always enabled.

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	5800	Users	3MB/XC	100		

Measurement Discussion: The following table shows that VM/ESA 2.3.0 has very similar performance characteristics relative to VM/ESA 2.2.0 for this workload and system configuration. External response time (AVG LAST(T)) improved by 2%, while the internal throughput rate (ITR(H)) decreased by 0.3%. Paging increased slightly due to a 1-page increase in the average working set size.

Table 1 (Page 1 of 3). Minidisk-only CMS-intensive migration from VM/ESA 2.2.0 on the 9121-742

Release Run ID	**2.2.0 S49E5800**	**2.3.0 S4AE5801**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5800**	**5800**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Response Time				
TRIV INT	0.033	0.033	0.000	0.00%
NONTRIV INT	0.248	0.249	0.001	0.40%
TOT INT	0.124	0.124	0.000	0.00%
TOT INT ADJ	0.192	0.193	0.001	0.52%
AVG FIRST (T)	0.293	0.292	-0.001	-0.50%
AVG LAST (T)	0.608	0.597	-0.012	-1.93%
Throughput				
AVG THINK (T)	26.57	26.51	-0.06	-0.22%
ETR	307.22	309.18	1.96	0.64%
ETR (T)	198.32	198.54	0.23	0.11%
ETR RATIO	1.549	1.557	0.008	0.52%
ITR (H)	222.85	222.30	-0.55	-0.25%
ITR	86.37	86.63	0.27	0.31%
EMUL ITR	137.20	137.47	0.27	0.20%
ITRR (H)	1.000	0.998	-0.002	-0.25%
ITRR	1.000	1.003	0.003	0.31%
Proc. Usage				
PBT/CMD (H)	17.949	17.994	0.044	0.25%
PBT/CMD	17.951	17.981	0.030	0.17%
CP/CMD (H)	7.078	7.067	-0.011	-0.15%
CP/CMD	6.656	6.648	-0.008	-0.11%
EMUL/CMD (H)	10.871	10.926	0.055	0.51%
EMUL/CMD	11.295	11.332	0.037	0.33%

Table 1 (Page 2 of 3). Minidisk-only CMS-intensive migration from VM/ESA 2.2.0 on the 9121-742

Release **Run ID**	**2.2.0** **S49E5800**	**2.3.0** **S4AE5801**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5800**	**5800**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
Processor Util.				
TOTAL (H)	355.97	357.26	1.29	0.36%
TOTAL	356.00	357.00	1.00	0.28%
UTIL/PROC (H)	88.99	89.31	0.32	0.36%
UTIL/PROC	89.00	89.25	0.25	0.28%
TOTAL EMUL (H)	215.59	216.94	1.34	0.62%
TOTAL EMUL	224.00	225.00	1.00	0.45%
MASTER TOTAL (H)	91.07	88.66	-2.41	-2.65%
MASTER TOTAL	91.00	89.00	-2.00	-2.20%
MASTER EMUL (H)	34.28	34.49	0.21	0.61%
MASTER EMUL	36.00	36.00	0.00	0.00%
TVR(H)	1.65	1.65	0.00	-0.26%
TVR	1.59	1.59	0.00	-0.16%
Storage				
NUCLEUS SIZE (V)	2844KB	2452KB	-392KB	-13.78%
TRACE TABLE (V)	650KB	650KB	0KB	0.00%
WKSET (V)	77	78	1	1.30%
PGBLPGS	230K	230K	0K	0.00%
PGBLPGS/USER	39.7	39.7	0.0	0.00%
TOT PAGES/USER (V)	191	193	2	1.05%
FREEPGS	17321	17555	234	1.35%
FREE UTIL	0.92	0.92	0.00	0.26%
SHRPGS	2158	2170	12	0.56%
Paging				
READS/SEC	1059	1107	48	4.53%
WRITES/SEC	797	834	37	4.64%
PAGE/CMD	9.359	9.776	0.417	4.46%
PAGE IO RATE (V)	341.000	334.400	-6.600	-1.94%
PAGE IO/CMD (V)	1.719	1.684	-0.035	-2.05%
XSTOR IN/SEC	555	529	-26	-4.68%
XSTOR OUT/SEC	1541	1562	21	1.36%
XSTOR/CMD	10.569	10.532	-0.037	-0.35%
FAST CLR/CMD	8.890	8.990	0.101	1.13%
Queues				
DISPATCH LIST	112.96	109.53	-3.44	-3.04%
ELIGIBLE LIST	0.02	0.04	0.02	100.00%
I/O				
VIO RATE	1793	1797	4	0.22%
VIO/CMD	9.041	9.051	0.010	0.11%
RIO RATE (V)	697	690	-7	-1.00%
RIO/CMD (V)	3.515	3.475	-0.039	-1.12%
NONPAGE RIO/CMD (V)	1.795	1.791	-0.004	-0.23%
DASD RESP TIME (V)	21.300	21.900	0.600	2.82%
MDC REAL SIZE (MB)	32.1	32.1	0.0	0.12%
MDC XSTOR SIZE (MB)	63.1	63.8	0.7	1.10%
MDC READS (I/Os)	569	572	3	0.53%
MDC WRITES (I/Os)	27	27	0	0.00%
MDC AVOID	529	533	4	0.76%
MDC HIT RATIO	0.92	0.93	0.01	1.09%

Table 1 (Page 3 of 3). Minidisk-only CMS-intensive migration from VM/ESA 2.2.0 on the 9121-742

Release **Run ID**	**2.2.0** **S49E5800**	**2.3.0** **S4AE5801**	**Difference**	**%Difference**
Environment				
Real Storage	**1024MB**	**1024MB**		
Exp. Storage	**1024MB**	**1024MB**		
Users	**5800**	**5800**		
VTAMs	**1**	**1**		
VSCSs	**3**	**3**		
Processors	**4**	**4**		
PRIVOPs				
PRIVOP/CMD	20.646	20.626	-0.020	-0.10%
DIAG/CMD	24.451	23.859	-0.591	-2.42%
DIAG 04/CMD	0.967	0.888	-0.078	-8.10%
DIAG 08/CMD	0.735	0.737	0.002	0.23%
DIAG 0C/CMD	0.212	0.192	-0.019	-9.08%
DIAG 14/CMD	0.024	0.025	0.000	0.35%
DIAG 58/CMD	1.249	1.250	0.001	0.07%
DIAG 98/CMD	0.211	0.209	-0.002	-1.18%
DIAG A4/CMD	3.558	3.582	0.024	0.67%
DIAG A8/CMD	2.678	2.683	0.005	0.18%
DIAG 214/CMD	12.522	11.976	-0.546	-4.36%
DIAG 270/CMD	0.921	0.941	0.020	2.14%
SIE/CMD	55.467	55.403	-0.064	-0.11%
SIE INTCPT/CMD	37.163	37.120	-0.043	-0.11%
FREE TOTL/CMD	45.326	44.892	-0.435	-0.96%
VTAM Machines				
WKSET (V)	4129	4140	11	0.27%
TOT CPU/CMD (V)	2.9498	2.9800	0.0302	1.02%
CP CPU/CMD (V)	1.3194	1.3263	0.0069	0.52%
VIRT CPU/CMD (V)	1.6304	1.6537	0.0233	1.43%
DIAG 98/CMD (V)	0.211	0.209	-0.002	-1.13%

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

9121-480 / Minidisk

Workload: FS8F0R

Hardware Configuration

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

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-3	RAMAC 2	4		8		16		
3390-2	3990-2	4	16		6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4			2			

Note: *R* next to the DASD counts means basic cache enabled; *W* means DASD fast write (and basic cache) enabled. RAMAC 2 refers to the RAMAC 2 Array Subsystem with 256MB cache and drawers in 3390-3 format. With the RAMAC 2 Array Subsystem, read and write caching are always enabled.

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	2040	Users	3MB/XC	100		

Measurement Discussion: External response time (AVG LAST(T)) increased by 1%, while internal throughput (ITR(H)) decreased by 0.5%. These results are similar to what was observed for the 9121-742 environment (see the previous section).

Table 2 (Page 1 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.2.0 on the 9121-480

Release **Run ID**	**2.2.0** **L29E2043**	**2.3.0** **L2AE2045**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2040**	**2040**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.131	0.135	0.004	3.05%
NONTRIV INT	0.437	0.451	0.014	3.20%
TOT INT	0.333	0.343	0.010	3.00%
TOT INT ADJ	0.298	0.308	0.009	3.12%
AVG FIRST (T)	0.270	0.269	-0.001	-0.55%
AVG LAST (T)	0.381	0.385	0.004	1.05%
Throughput				
AVG THINK (T)	26.18	26.14	-0.04	-0.13%
ETR	64.15	64.11	-0.04	-0.06%
ETR (T)	71.58	71.46	-0.12	-0.17%
ETR RATIO	0.896	0.897	0.001	0.11%
ITR (H)	79.26	78.90	-0.36	-0.45%
ITR	35.53	35.42	-0.11	-0.30%
EMUL ITR	53.00	52.97	-0.04	-0.07%
ITRR (H)	1.000	0.995	-0.005	-0.45%
ITRR	1.000	0.997	-0.003	-0.30%
Proc. Usage				
PBT/CMD (H)	25.234	25.349	0.115	0.46%
PBT/CMD	25.286	25.330	0.044	0.17%
CP/CMD (H)	8.918	8.980	0.062	0.69%
CP/CMD	8.382	8.397	0.015	0.17%
EMUL/CMD (H)	16.316	16.370	0.053	0.33%
EMUL/CMD	16.904	16.933	0.029	0.17%
Processor Util.				
TOTAL (H)	180.63	181.14	0.51	0.28%
TOTAL	181.00	181.00	0.00	0.00%
UTIL/PROC (H)	90.32	90.57	0.25	0.28%
UTIL/PROC	90.50	90.50	0.00	0.00%
TOTAL EMUL (H)	116.79	116.97	0.18	0.15%
TOTAL EMUL	121.00	121.00	0.00	0.00%
MASTER TOTAL (H)	90.05	90.20	0.16	0.17%
MASTER TOTAL	90.00	90.00	0.00	0.00%
MASTER EMUL (H)	51.58	51.55	-0.03	-0.06%
MASTER EMUL	54.00	54.00	0.00	0.00%
TVR(H)	1.55	1.55	0.00	0.13%
TVR	1.50	1.50	0.00	0.00%
Storage				
NUCLEUS SIZE (V)	2388KB	2452KB	64KB	2.68%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	82	83	1	1.22%
PGBLPGS	54133	54117	-16	-0.03%
PGBLPGS/USER	26.5	26.5	0.0	-0.03%
TOT PAGES/USER (V)	172	174	2	1.16%
FREEPGS	6196	6268	72	1.16%
FREE UTIL	0.95	0.94	-0.01	-1.15%
SHRPGS	1503	1488	-15	-1.00%

Table 2 (Page 2 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.2.0 on the 9121-480				
Release	**2.2.0**	**2.3.0**		
Run ID	**L29E2043**	**L2AE2045**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2040**	**2040**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	715	722	7	0.98%
WRITES/SEC	469	472	3	0.64%
PAGE/CMD	16.541	16.709	0.169	1.02%
PAGE IO RATE (V)	196.200	198.900	2.700	1.38%
PAGE IO/CMD (V)	2.741	2.783	0.043	1.55%
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	8.773	8.858	0.085	0.97%
Queues				
DISPATCH LIST	45.48	45.44	-0.04	-0.08%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	711	709	-2	-0.28%
VIO/CMD	9.933	9.922	-0.011	-0.11%
RIO RATE (V)	400	401	1	0.25%
RIO/CMD (V)	5.588	5.612	0.024	0.42%
NONPAGE RIO/CMD (V)	2.847	2.828	-0.019	-0.66%
DASD RESP TIME (V)	18.600	18.500	-0.100	-0.54%
MDC REAL SIZE (MB)	41.7	40.9	-0.8	-1.90%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	205	206	1	0.49%
MDC WRITES (I/Os)	9.80	9.91	0.11	1.12%
MDC AVOID	192	194	2	1.04%
MDC HIT RATIO	0.94	0.94	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.897	13.870	-0.027	-0.19%
DIAG/CMD	26.628	26.051	-0.577	-2.17%
DIAG 04/CMD	2.250	2.245	-0.005	-0.21%
DIAG 08/CMD	0.732	0.733	0.001	0.19%
DIAG 0C/CMD	0.212	0.192	-0.020	-9.29%
DIAG 14/CMD	0.024	0.024	0.000	0.39%
DIAG 58/CMD	1.251	1.250	-0.001	-0.07%
DIAG 98/CMD	1.134	1.104	-0.030	-2.68%
DIAG A4/CMD	3.569	3.593	0.024	0.66%
DIAG A8/CMD	2.663	2.666	0.003	0.12%
DIAG 214/CMD	12.495	11.927	-0.569	-4.55%
DIAG 270/CMD	0.921	0.941	0.020	2.18%
SIE/CMD	53.938	53.934	-0.004	-0.01%
SIE INTCPT/CMD	34.521	34.518	-0.003	-0.01%
FREE TOTL/CMD	49.733	49.484	-0.249	-0.50%
VTAM Machines				
WKSET (V)	560	573	13	2.32%
TOT CPU/CMD (V)	3.9737	3.9806	0.0069	0.17%
CP CPU/CMD (V)	1.4824	1.4927	0.0103	0.69%
VIRT CPU/CMD (V)	2.4913	2.4879	-0.0034	-0.14%
DIAG 98/CMD (V)	1.134	1.103	-0.031	-2.74%
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

9121-480 / SFS

Workload: FS8FMAXR

Hardware Configuration

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

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-3	RAMAC 2	4		8		16		
3390-2	3990-2	4	16		6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4			2			

Note: *R* next to the DASD counts means basic cache enabled; *W* means DASD fast write (and basic cache) enabled. RAMAC 2 refers to the RAMAC 2 Array Subsystem with 256MB cache and drawers in 3390-3 format. With the RAMAC 2 Array Subsystem, read and write caching are always enabled.

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 3%, while internal throughput (ITR(H)) decreased by 0.7%. The ITR decrease is somewhat larger than the 0.5% decrease observed for the 9121-480 minidisk environment (see the previous section) due to the fact that there was, additionally, some growth in SFS server CPU usage.

Table 3 (Page 1 of 3). SFS CMS-intensive migration from VM/ESA 2.2.0 on the 9121-480

Release Run ID	2.2.0 L29S1722	2.3.0 L2AS1721	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.127	0.130	0.003	2.36%
NONTRIV INT	0.474	0.465	-0.009	-1.90%
TOT INT	0.357	0.354	-0.003	-0.84%
TOT INT ADJ	0.316	0.312	-0.004	-1.23%
AVG FIRST (T)	0.255	0.251	-0.004	-1.76%
AVG LAST (T)	0.384	0.373	-0.011	-2.74%
Throughput				
AVG THINK (T)	26.16	26.20	0.04	0.13%
ETR	53.24	53.01	-0.23	-0.43%
ETR (T)	60.22	60.20	-0.02	-0.04%
ETR RATIO	0.884	0.881	-0.004	-0.40%
ITR (H)	66.79	66.32	-0.46	-0.69%
ITR	29.54	29.22	-0.32	-1.10%
EMUL ITR	44.82	44.44	-0.38	-0.84%
ITRR (H)	1.000	0.993	-0.007	-0.69%
ITRR	1.000	0.989	-0.011	-1.10%
Proc. Usage				
PBT/CMD (H)	29.947	30.156	0.209	0.70%
PBT/CMD	29.889	30.232	0.343	1.15%
CP/CMD (H)	10.872	10.993	0.121	1.11%
CP/CMD	10.129	10.465	0.336	3.32%
EMUL/CMD (H)	19.074	19.162	0.088	0.46%
EMUL/CMD	19.760	19.767	0.007	0.04%
Processor Util.				
TOTAL (H)	180.35	181.54	1.20	0.66%
TOTAL	180.00	182.00	2.00	1.11%
UTIL/PROC (H)	90.17	90.77	0.60	0.66%
UTIL/PROC	90.00	91.00	1.00	1.11%
TOTAL EMUL (H)	114.87	115.36	0.49	0.43%
TOTAL EMUL	119.00	119.00	0.00	0.00%
MASTER TOTAL (H)	89.97	90.55	0.58	0.64%
MASTER TOTAL	90.00	91.00	1.00	1.11%
MASTER EMUL (H)	51.82	51.90	0.08	0.16%
MASTER EMUL	54.00	54.00	0.00	0.00%
TVR(H)	1.57	1.57	0.00	0.24%
TVR	1.51	1.53	0.02	1.11%
Storage				
NUCLEUS SIZE (V)	2388KB	2452KB	64KB	2.68%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	79	81	2	2.53%
PGBLPGS	55319	55389	70	0.13%
PGBLPGS/USER	32.2	32.2	0.0	0.13%
TOT PAGES/USER (V)	157	159	2	1.27%
FREEPGS	5385	5377	-8	-0.15%
FREE UTIL	0.95	0.95	0.00	0.15%
SHRPGS	1706	1718	12	0.70%

Table 3 (Page 2 of 3). SFS CMS-intensive migration from VM/ESA 2.2.0 on the 9121-480

Release Run ID	2.2.0 L29S1722	2.3.0 L2AS1721	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	588	605	17	2.89%
WRITES/SEC	388	397	9	2.32%
PAGE/CMD	16.206	16.644	0.438	2.70%
PAGE IO RATE (V)	148.100	154.600	6.500	4.39%
PAGE IO/CMD (V)	2.459	2.568	0.109	4.43%
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	8.518	8.654	0.136	1.60%
Queues				
DISPATCH LIST	41.20	41.14	-0.05	-0.13%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	610	609	-1	-0.16%
VIO/CMD	10.129	10.116	-0.013	-0.13%
RIO RATE (V)	362	367	5	1.38%
RIO/CMD (V)	6.011	6.096	0.085	1.42%
NONPAGE RIO/CMD (V)	3.552	3.528	-0.024	-0.67%
DASD RESP TIME (V)	17.000	17.000	0.000	0.00%
MDC REAL SIZE (MB)	69.5	67.0	-2.5	-3.60%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	161	162	1	0.62%
MDC WRITES (I/Os)	15	15	0	0.00%
MDC AVOID	134	135	1	0.75%
MDC HIT RATIO	0.83	0.83	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	20.501	20.673	0.172	0.84%
DIAG/CMD	24.847	24.351	-0.497	-2.00%
DIAG 04/CMD	2.489	2.488	-0.002	-0.06%
DIAG 08/CMD	0.736	0.736	0.000	0.05%
DIAG 0C/CMD	0.237	0.217	-0.020	-8.50%
DIAG 14/CMD	0.025	0.025	0.000	0.34%
DIAG 58/CMD	1.250	1.249	0.000	-0.03%
DIAG 98/CMD	1.387	1.347	-0.040	-2.87%
DIAG A4/CMD	1.985	1.992	0.007	0.38%
DIAG A8/CMD	2.456	2.457	0.001	0.05%
DIAG 214/CMD	11.992	11.522	-0.470	-3.92%
DIAG 270/CMD	0.920	0.941	0.021	2.26%
SIE/CMD	61.787	62.141	0.354	0.57%
SIE INTCPT/CMD	42.015	42.256	0.241	0.57%
FREE TOTL/CMD	53.169	53.171	0.002	0.00%
VTAM Machines				
WKSET (V)	509	495	-14	-2.75%
TOT CPU/CMD (V)	4.2250	4.2173	-0.0077	-0.18%
CP CPU/CMD (V)	1.5775	1.5688	-0.0087	-0.55%
VIRT CPU/CMD (V)	2.6476	2.6485	0.0009	0.03%
DIAG 98/CMD (V)	1.387	1.347	-0.040	-2.86%

Table 3 (Page 3 of 3). SFS CMS-intensive migration from VM/ESA 2.2.0 on the 9121-480

Release **Run ID**	**2.2.0** **L29S1722**	**2.3.0** **L2AS1721**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
SFS Servers				
WKSET (V)	3364	3360	-4	-0.12%
TOT CPU/CMD (V)	3.4040	3.4698	0.0658	1.93%
CP CPU/CMD (V)	1.4944	1.5319	0.0375	2.51%
VIRT CPU/CMD (V)	1.9096	1.9379	0.0283	1.48%
FP REQ/CMD(Q)	1.114	1.151	0.037	3.32%
IO/CMD (Q)	1.614	1.644	0.030	1.86%
IO TIME/CMD (Q)	0.023	0.023	0.000	0.00%
SFS TIME/CMD (Q)	0.035	0.034	-0.001	-2.86%

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 4 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 4. SFS CMS-intensive migration from VM/ESA 2.2.0 on the 9121-480

Release Run ID	2.2.0 L29S1722	2.3.0 L2AS1721	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Close File Requests	0.3619	0.3796	0.0177	4.89%
Commit Requests	0.0163	0.0164	0.0001	0.61%
Connect Requests	0.0078	0.0079	0.0001	1.28%
Delete File Requests	0.0734	0.0736	0.0002	0.27%
Lock Requests	0.0246	0.0246	0.0000	0.00%
Open File New Requests	0.0033	0.0033	0.0000	0.00%
Open File Read Requests	0.2164	0.2341	0.0177	8.18%
Open File Replace Requests	0.1210	0.1210	0.0000	0.00%
Open File Write Requests	0.0212	0.0212	0.0000	0.00%
Query File Pool Requests	0.0000	0.0000	0.0000	na
Query User Space Requests	0.0212	0.0213	0.0001	0.47%
Read File Requests	0.1441	0.1442	0.0001	0.07%
Refresh Directory Requests	0.0227	0.0229	0.0002	0.88%
Rename Requests	0.0049	0.0049	0.0000	0.00%
Unlock Requests	0.0246	0.0246	0.0000	0.00%
Write File Requests	0.0507	0.0509	0.0002	0.39%
Total File Pool Requests	1.1143	1.1506	0.0363	3.26%
File Pool Request Service Time	35.1515	34.2354	-0.9161	-2.61%
Local File Pool Requests	1.1143	1.1506	0.0363	3.26%
Begin LUWs	0.4451	0.4633	0.0182	4.09%
Agent Holding Time (msec)	110.7923	106.4140	-4.3783	-3.95%
SAC Calls	5.4834	5.6300	0.1466	2.67%
Catalog Lock Conflicts	0.0018	0.0012	-0.0006	-33.33%
Total Lock Conflicts	0.0018	0.0012	-0.0006	-33.33%
Lock Wait Time (msec)	0.1426	0.0626	-0.0800	-56.10%
File Blocks Read	0.8995	0.9175	0.0180	2.00%
File Blocks Written	0.4966	0.4984	0.0018	0.36%
Catalog Blocks Read	0.5195	0.5269	0.0074	1.42%
Catalog Blocks Written	0.2696	0.2733	0.0037	1.37%
Control Minidisk Blocks Written	0.0511	0.0527	0.0016	3.13%
Log Blocks Written	0.4615	0.4636	0.0021	0.46%
Total DASD Block Transfers	2.6976	2.7324	0.0348	1.29%
BIO Requests to Read File Block	0.3880	0.4057	0.0177	4.56%
BIO Requests to Write File Blocks	0.1790	0.1792	0.0002	0.11%
BIO Requests to Read Catalog Blks	0.5195	0.5269	0.0074	1.42%
BIO Requests to Write Catalog Blks	0.2231	0.2239	0.0008	0.36%
BIO Requests to Write Ctl Mdsk Blks	0.0021	0.0021	0.0000	0.00%
BIO Requests to Write Log Blocks	0.4026	0.4046	0.0020	0.50%
Total BIO Requests	1.7142	1.7425	0.0283	1.65%
Total BIO Request Time (msec)	23.1710	22.8600	-0.3110	-1.34%
I/O Requests to Read File Blocks	0.2637	0.2818	0.0181	6.86%
I/O Requests to Write File Blocks	0.1947	0.1956	0.0009	0.46%
I/O Requests to Read Catalog Blks	0.5195	0.5269	0.0074	1.42%
I/O Requests to Write Catalog Blks	0.2293	0.2307	0.0014	0.61%
I/O Requests to Write Ctl Mdsk Blks	0.0038	0.0040	0.0002	5.26%
I/O Requests to Write Log Blocks	0.4029	0.4050	0.0021	0.52%
Total I/O Requests	1.6138	1.6439	0.0301	1.87%
Get Logname Requests	0.0032	0.0033	0.0001	3.13%
Get LUWID Requests	0.0032	0.0033	0.0001	3.13%
Total CRR Requests	0.0065	0.0066	0.0001	1.54%
CRR Request Service Time (msec)	0.0495	0.0430	-0.0065	-13.13%
Log I/O Requests	0.0065	0.0066	0.0001	1.54%

Note: Query Filepool Counters — normalized by command

Table 5 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 5. SFS CMS-intensive migration from VM/ESA 2.2.0 on the 9121-480

Release **Run ID**	**2.2.0** **L29S1722**	**2.3.0** **L2AS1721**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1720**	**1720**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Agents Held	6.7	6.4	-0.3	-3.99%
Agents In-call	2.1	2.1	-0.1	-2.64%
Avg LUW Time (msec)	248.9	229.7	-19.2	-7.72%
Avg File Pool Request Time (msec)	31.5	29.8	-1.8	-5.68%
Avg Lock Wait Time (msec)	79.2	52.2	-27.1	-34.15%
SAC Calls / FP Request	4.92	4.89	-0.03	-0.57%
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)	853	840	-13	-1.52%
Checkpoints Taken (delta)	34	35	1	2.94%
Checkpoint Duration (sec)	2.8	2.7	-0.1	-3.50%
Seconds Between Checkpoints	56.5	54.9	-1.6	-2.75%
Checkpoint Utilization	5.0	5.0	0.0	-0.63%
BIO Request Time (msec)	13.52	13.12	-0.40	-2.94%
Blocking Factor (Blocks/BIO)	1.57	1.57	-0.01	-0.36%
Chaining Factor (Blocks/IO)	1.67	1.66	-0.01	-0.56%

Note: Query Filepool Counters — derived results

Table 6 compares the VM/ESA 2.3.0 SFS measurement to the corresponding VM/ESA 2.3.0 minidisk-only measurement from Table 2 on page 18.

Table 6 (Page 1 of 3). Minidisk to SFS comparison for VM/ESA 2.3.0 on the 9121-480

File System Release Run ID	Minidisk 2.3.0 L2AE2045	SFS 2.3.0 L2AS1721	Difference	%Difference
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2040**	**1720**	**-320**	**-15.69%**
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.135	0.130	-0.005	-3.70%
NONTRIV INT	0.451	0.465	0.014	3.10%
TOT INT	0.343	0.354	0.011	3.21%
TOT INT ADJ	0.308	0.312	0.004	1.29%
AVG FIRST (T)	0.269	0.251	-0.018	-6.88%
AVG LAST (T)	0.385	0.373	-0.012	-3.12%
Throughput				
AVG THINK (T)	26.14	26.20	0.06	0.21%
ETR	64.11	53.01	-11.10	-17.31%
ETR (T)	71.46	60.20	-11.26	-15.75%
ETR RATIO	0.897	0.881	-0.017	-1.86%
ITR (H)	78.90	66.32	-12.58	-15.94%
ITR	35.42	29.22	-6.20	-17.50%
EMUL ITR	52.97	44.44	-8.53	-16.10%
ITRR (H)	1.000	0.841	-0.159	-15.94%
ITRR	1.000	0.825	-0.175	-17.50%
Proc. Usage				
PBT/CMD (H)	25.349	30.156	4.807	18.96%
PBT/CMD	25.330	30.232	4.902	19.35%
CP/CMD (H)	8.980	10.993	2.014	22.42%
CP/CMD	8.397	10.465	2.068	24.63%
EMUL/CMD (H)	16.370	19.162	2.793	17.06%
EMUL/CMD	16.933	19.767	2.834	16.73%
Processor Util.				
TOTAL (H)	181.14	181.54	0.41	0.22%
TOTAL	181.00	182.00	1.00	0.55%
UTIL/PROC (H)	90.57	90.77	0.20	0.22%
UTIL/PROC	90.50	91.00	0.50	0.55%
TOTAL EMUL (H)	116.97	115.36	-1.61	-1.38%
TOTAL EMUL	121.00	119.00	-2.00	-1.65%
MASTER TOTAL (H)	90.20	90.55	0.35	0.38%
MASTER TOTAL	90.00	91.00	1.00	1.11%
MASTER EMUL (H)	51.55	51.90	0.36	0.69%
MASTER EMUL	54.00	54.00	0.00	0.00%
TVR(H)	1.55	1.57	0.03	1.62%
TVR	1.50	1.53	0.03	2.24%
Storage				
NUCLEUS SIZE (V)	2452KB	2452KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	83	81	-2	-2.41%
PGBLPGS	54117	55389	1272	2.35%
PGBLPGS/USER	26.5	32.2	5.7	21.39%
TOT PAGES/USER (V)	174	159	-15	-8.62%
FREEPGS	6268	5377	-891	-14.22%
FREE UTIL	0.94	0.95	0.01	1.37%
SHRPGS	1488	1718	230	15.46%

Table 6 (Page 2 of 3). Minidisk to SFS comparison for VM/ESA 2.3.0 on the 9121-480				
File System **Release** **Run ID**	**Minidisk** **2.3.0** **L2AE2045**	**SFS** **2.3.0** **L2AS1721**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2040**	**1720**	**-320**	**-15.69%**
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	722	605	-117	-16.20%
WRITES/SEC	472	397	-75	-15.89%
PAGE/CMD	16.709	16.644	-0.065	-0.39%
PAGE IO RATE (V)	198.900	154.600	-44.300	-22.27%
PAGE IO/CMD (V)	2.783	2.568	-0.215	-7.74%
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	8.858	8.654	-0.204	-2.31%
Queues				
DISPATCH LIST	45.44	41.14	-4.30	-9.47%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	709	609	-100	-14.10%
VIO/CMD	9.922	10.116	0.194	1.95%
RIO RATE (V)	401	367	-34	-8.48%
RIO/CMD (V)	5.612	6.096	0.484	8.63%
NONPAGE RIO/CMD (V)	2.828	3.528	0.700	24.75%
DASD RESP TIME (V)	18.500	17.000	-1.500	-8.11%
MDC REAL SIZE (MB)	40.9	67.0	26.1	63.86%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	206	162	-44	-21.36%
MDC WRITES (I/Os)	9.91	15	5.09	51.36%
MDC AVOID	194	135	-59	-30.41%
MDC HIT RATIO	0.94	0.83	-0.11	-11.70%
PRIVOPs				
PRIVOP/CMD	13.870	20.673	6.803	49.04%
DIAG/CMD	26.051	24.351	-1.700	-6.53%
DIAG 04/CMD	2.245	2.488	0.242	10.79%
DIAG 08/CMD	0.733	0.736	0.003	0.38%
DIAG 0C/CMD	0.192	0.217	0.025	12.86%
DIAG 14/CMD	0.024	0.025	0.000	0.89%
DIAG 58/CMD	1.250	1.249	0.000	-0.03%
DIAG 98/CMD	1.104	1.347	0.243	22.05%
DIAG A4/CMD	3.593	1.992	-1.601	-44.55%
DIAG A8/CMD	2.666	2.457	-0.210	-7.86%
DIAG 214/CMD	11.927	11.522	-0.404	-3.39%
DIAG 270/CMD	0.941	0.941	0.000	0.02%
SIE/CMD	53.934	62.141	8.206	15.22%
SIE INTCPT/CMD	34.518	42.256	7.738	22.42%
FREE TOTL/CMD	49.484	53.171	3.687	7.45%
VTAM Machines				
WKSET (V)	573	495	-78	-13.61%
TOT CPU/CMD (V)	3.9806	4.2173	0.2367	5.95%
CP CPU/CMD (V)	1.4927	1.5688	0.0761	5.10%
VIRT CPU/CMD (V)	2.4879	2.6485	0.1606	6.46%
DIAG 98/CMD (V)	1.103	1.347	0.244	22.16%

Table 6 (Page 3 of 3). Minidisk to SFS comparison for VM/ESA 2.3.0 on the 9121-480				
File System **Release** **Run ID**	**Minidisk** **2.3.0** **L2AE2045**	**SFS** **2.3.0** **L2AS1721**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2040**	**1720**	**-320**	**-15.69%**
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
SFS Servers				
WKSET (V)	na	3360		
TOT CPU/CMD (V)	na	3.4698		
CP CPU/CMD (V)	na	1.5319		
VIRT CPU/CMD (V)	na	1.9379		
FP REQ/CMD(Q)	na	1.151		
IO/CMD (Q)	na	1.644		
IO TIME/CMD (Q)	na	0.023		
SFS TIME/CMD (Q)	na	0.034		
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.

VSE/ESA Guest

This section examines the performance impact of migrating a VSE/ESA 2.1.0 guest from VM/ESA 2.2.0 to VM/ESA 2.3.0 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 78 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*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*VSAM*	*VSE Sys.*	*VM Sys.*
3380-A	3880-03	2						1
3390-2	3990-02	4				10	2	
3380-K	3990-03	4				10		

Software Configuration

VSE version: 2.1.0 (using the standard dispatcher)

Virtual Machines:

Virtual Machine	*Number*	*Type*	*Machine Size/Mode*	*SHARE*	*RESERVED*	*Other Options*
VSEVR	1	VSE V=R	96MB/ESA	100		IOASSIST ON CCWTRANS OFF
or						
VSEVV	1	VSE V=V	96MB/ESA	100		IOASSIST OFF
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: The measured internal throughput rate (ITR(H)) decreased somewhat relative to VM/ESA 2.2.0. The amount of decrease was 0.8% for the V=R environment and 1.0% for the V=V environment. Elapsed times were equivalent within measurement variability.

Table 7 (Page 1 of 2). VSE/ESA V=R guest migration from VM/ESA 2.2.0 on the 9121-320

Release **Run ID**	**2.2.0** **L1R98PF2**	**2.3.0** **L1RA8PF2**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	874.0	871.0	-3.0	-0.34%
ETR (C)	7.69	7.72	0.03	0.34%
ITR (H)	18.98	18.83	-0.15	-0.79%
ITR	18.75	18.82	0.06	0.34%
ITRR (H)	1.000	0.992	-0.008	-0.79%
ITRR	1.000	1.003	0.003	0.34%
Proc. Usage (Sec)				
PBT/CMD (H)	3.162	3.187	0.025	0.80%
PBT/CMD	3.199	3.188	-0.011	-0.34%
CP/CMD (H)	0.276	0.291	0.014	5.14%
CP/CMD	0.312	0.233	-0.079	-25.26%
EMUL/CMD (H)	2.885	2.896	0.011	0.38%
EMUL/CMD	2.887	2.955	0.068	2.35%
Processor Util.				
TOTAL (H)	40.51	40.98	0.46	1.15%
TOTAL	41.00	41.00	0.00	0.00%
TOTAL EMUL (H)	36.97	37.24	0.27	0.73%
TOTAL EMUL	37.00	38.00	1.00	2.70%
TVR(H)	1.10	1.10	0.00	0.41%
TVR	1.11	1.08	-0.03	-2.63%
Storage				
NUCLEUS SIZE (V)	2820KB	2884KB	64KB	2.27%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	38541	38532	-9	-0.02%
FREEPGS	89	85	-4	-4.49%
FREE UTIL	0.56	0.61	0.05	9.42%
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	7.804	7.777	-0.027	-0.34%
RIO RATE (V)	2.000	2.000	0.000	0.00%
RIO/CMD (V)	15.607	15.554	-0.054	-0.34%
DASD IO TOTAL (V)	350630	350840	210	0.06%
DASD IO RATE (V)	389.59	389.82	0.23	0.06%
DASD IO/CMD (V)	3040.18	3031.56	-8.62	-0.28%
MDC REAL SIZE (MB)	8.1	8.0	-0.1	-0.97%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	0.03	0.03	0	0.00%
MDC WRITES (I/Os)	0.03	0.03	0	0.00%
MDC AVOID	0.01	0.01	0	0.00%
MDC HIT RATIO	0.30	0.32	0.02	6.67%

Table 7 (Page 2 of 2). VSE/ESA V=R guest migration from VM/ESA 2.2.0 on the 9121-320				
Release **Run ID**	**2.2.0** **L1R98PF2**	**2.3.0** **L1RA8PF2**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD (R)	10.791	10.675	-0.116	-1.07%
DIAG/CMD (R)	606.378	607.677	1.300	0.21%
SIE/CMD	2739.054	2784.089	45.036	1.64%
SIE INTCPT/CMD	2273.414	2310.794	37.380	1.64%
FREE TOTL/CMD	546.250	528.821	-17.429	-3.19%
Note: V=VMPRF, H=Hardware Monitor, C=VSE console, Unmarked=RTM				

Table 8 (Page 1 of 2). VSE/ESA V=V guest migration from VM/ESA 2.2.0 on the 9121-320				
Release **Run ID**	**2.2.0** **L1V98PF2**	**2.3.0** **L1VA8PF2**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (C)	484.0	483.0	-1.0	-0.21%
ETR (C)	13.88	13.91	0.03	0.21%
ITR (H)	14.76	14.61	-0.15	-1.03%
ITR	14.77	14.65	-0.13	-0.85%
ITRR (H)	1.000	0.990	-0.010	-1.03%
ITRR	1.000	0.992	-0.008	-0.85%
Proc. Usage (Sec)				
PBT/CMD (H)	4.065	4.107	0.042	1.04%
PBT/CMD	4.062	4.097	0.035	0.86%
CP/CMD (H)	1.124	1.157	0.033	2.92%
CP/CMD	0.994	1.035	0.041	4.13%
EMUL/CMD (H)	2.941	2.951	0.010	0.32%
EMUL/CMD	3.068	3.062	-0.006	-0.21%
Processor Util.				
TOTAL (H)	94.06	95.24	1.18	1.25%
TOTAL	94.00	95.00	1.00	1.06%
TOTAL EMUL (H)	68.06	68.42	0.36	0.53%
TOTAL EMUL	71.00	71.00	0.00	0.00%
TVR(H)	1.38	1.39	0.01	0.72%
TVR	1.32	1.34	0.01	1.06%

Table 8 (Page 2 of 2). VSE/ESA V=V guest migration from VM/ESA 2.2.0 on the 9121-320

Release **Run ID**	**2.2.0** **L1V98PF2**	**2.3.0** **L1VA8PF2**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=R**	**V=R**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Storage				
NUCLEUS SIZE (V)	2820KB	2884KB	64KB	2.27%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	38458	38459	1	0.00%
FREEPGS	108	107	-1	-0.93%
FREE UTIL	0.62	0.62	0.00	0.18%
Paging				
PAGE/CMD	47.536	112.125	64.589	135.88%
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	280.893	276.000	-4.893	-1.74%
I/O				
VIO RATE	724.000	722.000	-2.000	-0.28%
VIO/CMD	3128.714	3113.625	-15.089	-0.48%
RIO RATE (V)	316.000	322.000	6.000	1.90%
RIO/CMD (V)	1365.571	1388.625	23.054	1.69%
DASD IO TOTAL (V)	151261	153866	2605	1.72%
DASD IO RATE (V)	315.13	320.55	5.43	1.72%
DASD IO/CMD (V)	1361.80	1382.39	20.59	1.51%
MDC REAL SIZE (MB)	112.4	112.6	0.2	0.19%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	443	442	-1	-0.23%
MDC WRITES (I/Os)	219	218	-1	-0.46%
MDC AVOID	415	407	-8	-1.93%
MDC HIT RATIO	0.86	0.85	-0.01	-1.16%
PRIVOPs				
PRIVOP/CMD (R)	3126.116	3117.492	-8.624	-0.28%
DIAG/CMD (R)	452.670	450.594	-2.076	-0.46%
SIE/CMD	13612.500	13618.875	6.375	0.05%
SIE INTCPT/CMD	11842.875	11848.421	5.546	0.05%
FREE TOTL/CMD	3980.036	4019.250	39.214	0.99%

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

TCP/IP

Telnet and FTP measurements were obtained to evaluate the performance effects of migrating from TCP/IP 2.4 to TCP/IP 310. The results show a 3.7% decrease in TCPIP virtual machine processor usage. This improvement resulted from pathlength improvements that were made to the TCP/IP stack.

Telnet

Workload: FS8F0R

Hardware Configuration

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

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-3	RAMAC 2	4		8		16		
3390-2	3990-2	4	16		6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4			2			

Note: *R* next to the DASD counts means basic cache enabled; *W* means DASD fast write (and basic cache) enabled. RAMAC 2 refers to the RAMAC 2 Array Subsystem with 256MB cache and drawers in 3390-3 format. With the RAMAC 2 Array Subsystem, read and write caching are always enabled.

Communications:

16 Mbit IBM Token Ring
3172-3 Interconnect Controller

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*
TCPIP	1	TCP/IP	256MB/XA	10000	2700	QUICKDSP ON
SMART	1	RTM	32MB/XA	3%	400	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1800	Users	3MB/XC	100		

Measurement Discussion: Relative to TCP/IP 2.4, external response time (AVG LAST(T)) improved by 8% while internal throughput rate (ITR(H)) improved by 0.6%. These results are due to reduced processor usage in the TCPIP virtual machine, as shown by TCPIP VIRT CPU/CMD(V) in Table 9. This and other performance improvements are described in "Performance Improvements" on page 3.

These measurements were run with the size of each TCP/IP buffer pool set to be slightly larger than maximum requirements, as determined by the NETSTAT POOLSIZE command. This is the "Trimmed" configuration in Table 21 on page 59. Because of this, these results do not reflect the additional benefits from the "Reduced TCP/IP Real Storage Usage" performance improvement. That improvement is discussed in "Reduced TCP/IP Real Storage Requirements" on page 58.

Table 9 (Page 1 of 3). Telnet migration on a 9121-480

TCP/IP Release **VM/ESA Release** **Run ID**	**2.4** **2.3.0** **L2AE1803**	**310** **2.3.0** **L2AE1804**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1800**	**1800**		
VTAMs	**na**	**na**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.122	0.111	-0.011	-9.02%
NONTRIV INT	0.837	0.812	-0.025	-2.99%
TOT INT	0.245	0.232	-0.013	-5.31%
TOT INT ADJ	0.227	0.214	-0.013	-5.60%
AVG FIRST (T)	0.331	0.299	-0.033	-9.95%
AVG LAST (T)	0.438	0.402	-0.036	-8.32%
Throughput				
AVG THINK (T)	24.53	24.50	-0.03	-0.12%
ETR	58.28	58.19	-0.09	-0.15%
ETR (T)	62.87	62.97	0.10	0.16%
ETR RATIO	0.927	0.924	-0.003	-0.31%
ITR (H)	69.84	70.24	0.40	0.57%
ITR	32.41	32.47	0.06	0.19%
EMUL ITR	52.76	53.28	0.52	0.98%
ITRR (H)	1.000	1.006	0.006	0.57%
ITRR	1.000	1.002	0.002	0.19%
Proc. Usage				
PBT/CMD (H)	28.635	28.474	-0.162	-0.56%
PBT/CMD	28.629	28.425	-0.204	-0.71%
CP/CMD (H)	11.688	11.767	0.079	0.67%
CP/CMD	10.975	11.116	0.141	1.29%
EMUL/CMD (H)	16.947	16.707	-0.240	-1.42%
EMUL/CMD	17.655	17.309	-0.346	-1.96%

Table 9 (Page 2 of 3). Telnet migration on a 9121-480

TCP/IP Release **VM/ESA Release** **Run ID**	**2.4** **2.3.0** **L2AE1803**	**310** **2.3.0** **L2AE1804**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1800**	**1800**		
VTAMs	**na**	**na**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Processor Util.				
TOTAL (H)	180.04	179.31	-0.73	-0.41%
TOTAL	180.00	179.00	-1.00	-0.56%
UTIL/PROC (H)	90.02	89.65	-0.37	-0.41%
UTIL/PROC	90.00	89.50	-0.50	-0.56%
TOTAL EMUL (H)	106.55	105.21	-1.34	-1.26%
TOTAL EMUL	111.00	109.00	-2.00	-1.80%
MASTER TOTAL (H)	90.57	90.28	-0.28	-0.31%
MASTER TOTAL	91.00	90.00	-1.00	-1.10%
MASTER EMUL (H)	44.04	43.42	-0.62	-1.41%
MASTER EMUL	46.00	45.00	-1.00	-2.17%
TVR(H)	1.69	1.70	0.01	0.86%
TVR	1.62	1.64	0.02	1.27%
Storage				
NUCLEUS SIZE (V)	2452KB	2452KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	86	85	-1	-1.16%
PGBLPGS	55511	55510	-1	0.00%
PGBLPGS/USER	30.8	30.8	0.0	0.00%
TOT PAGES/USER (V)	179	179	0	0.00%
FREEPGS	5349	5355	6	0.11%
FREE UTIL	0.96	0.96	0.00	-0.11%
SHRPGS	917	961	44	4.80%
Paging				
READS/SEC	690	692	2	0.29%
WRITES/SEC	410	414	4	0.98%
PAGE/CMD	17.496	17.563	0.067	0.39%
PAGE IO RATE (V)	170.900	171.700	0.800	0.47%
PAGE IO/CMD (V)	2.718	2.727	0.008	0.31%
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	8.891	8.909	0.018	0.20%
Queues				
DISPATCH LIST	23.40	22.33	-1.07	-4.57%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	880	884	4	0.45%
VIO/CMD	13.996	14.038	0.041	0.29%
RIO RATE (V)	615	618	3	0.49%
RIO/CMD (V)	9.782	9.814	0.032	0.33%
NONPAGE RIO/CMD (V)	7.063	7.087	0.024	0.34%
DASD RESP TIME (V)	18.100	18.100	0.000	0.00%
MDC REAL SIZE (MB)	42.4	42.0	-0.4	-1.00%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	181	181	0	0.00%
MDC WRITES (I/Os)	8.61	8.63	0.02	0.23%
MDC AVOID	171	171	0	0.00%
MDC HIT RATIO	0.94	0.94	0.00	0.00%

Table 9 (Page 3 of 3). Telnet migration on a 9121-480				
TCP/IP Release **VM/ESA Release** **Run ID**	**2.4** **2.3.0** **L2AE1803**	**310** **2.3.0** **L2AE1804**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1800**	**1800**		
VTAMs	**na**	**na**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
PRIVOPs				
PRIVOP/CMD	1.618	1.626	0.009	0.54%
DIAG/CMD	36.101	36.180	0.079	0.22%
DIAG 04/CMD	2.499	2.496	-0.003	-0.10%
DIAG 08/CMD	0.732	0.736	0.004	0.57%
DIAG 0C/CMD	0.192	0.192	0.000	0.26%
DIAG 14/CMD	0.024	0.025	0.000	0.95%
DIAG 58/CMD	1.248	1.250	0.002	0.17%
DIAG 98/CMD	5.161	5.190	0.029	0.57%
DIAG A4/CMD	3.593	3.591	-0.001	-0.04%
DIAG A8/CMD	2.666	2.685	0.018	0.69%
DIAG 214/CMD	11.922	11.910	-0.012	-0.10%
DIAG 270/CMD	0.940	0.940	0.000	0.02%
SIE/CMD	65.593	65.726	0.134	0.20%
SIE INTCPT/CMD	43.947	44.694	0.747	1.70%
FREE TOTL/CMD	53.600	53.705	0.105	0.20%
TCPIP Machine				
WKSET (V)	2700	2700	0	0.00%
TOT CPU/CMD (V)	5.5580	5.3550	-0.2030	-3.65%
CP CPU/CMD (V)	2.4030	2.4260	0.0230	0.96%
VIRT CPU/CMD (V)	3.1550	2.9290	-0.2260	-7.16%
DIAG 98/CMD (V)	5.164	5.190	0.026	0.50%
Note: T=TPNS, V=VMPRF, H=Hardware Monitor,Unmarked=RTM				

FTP

Measurement Description: The measured system consisted of a 9121-480 (2-way) processor with 256MB of central storage and no expanded storage. It was running VM/ESA 2.3.0 with TCP/IP 2.4 or TCP/IP 310. This system was attached to a 16Mbit IBM Token Ring through a 3172-3 Interconnect Controller. The TCPIP server was run with the “trimmed” buffer pool sizes listed in Table 21 on page 59 and with a packet (MTU) size of 4000 bytes. File transfer was to/from an RS/6000 model 520 on the same token ring. The RS/6000 was running AIX 3.2.4.

The workload consisted of “n” consecutive identical FTP file transfers (get or put) initiated by a client virtual machine on the VM/ESA system. These file transfers were to/from the client virtual machine′s 191 minidisk, which was defined in RAMAC 2 Array Subsystem storage. Minidisk cache was enabled for this minidisk.

Measurement Results: The measurement results are summarized in Table 10, Table 11, and Table 12.

Table 10. FTP Performance on a 9121-480: Get 2MB Files

Run Id	**G2M_A243**	**G2M_A313**	**G2M_B243**	**G2M_B313**
Direction	**Get**	**Get**	**Get**	**Get**
Mode	**ASCII**	**ASCII**	**Binary**	**Binary**
TCP/IP Release	**2.4**	**310**	**2.4**	**310**
File Size (Kb)	2073.8	2073.8	2073.8	2073.8
Files Transferred	5	5	5	5
Total Kb	10369	10369	10369	10369
Elapsed Time (sec)	42.0	42.0	42.1	42.4
Rate (Kb/sec)	246.9	246.9	246.3	244.6
Rate Ratio	1.000	1.000	1.000	0.993
TCPIP CPU Usage				
CP CPU (sec)	1.48	1.50	1.51	1.51
Virtual CPU (sec)	2.52	2.34	2.54	2.33
Total CPU (sec)	4.00	3.84	4.05	3.84
Total CPU Ratio	1.000	0.960	1.000	0.948
Total CPU/Kb (msec)	0.386	0.370	0.391	0.370
VM Client CPU Usage				
CP CPU (sec)	0.56	0.54	0.53	0.52
Virtual CPU (sec)	10.19	5.10	0.85	0.81
Total CPU (sec)	10.75	5.64	1.38	1.33
Total CPU Ratio	1.000	0.525	1.000	0.964
Total CPU/Kb (msec)	1.037	0.544	0.133	0.128
Total CPU Usage				
Total CPU (sec)	14.75	9.48	5.43	5.17
Total CPU Ratio	1.000	0.643	1.000	0.952
Total CPU/Kb (msec)	1.423	0.914	0.524	0.499

Table 11. FTP Performance on a 9121-480: Get 24KB Files

Run Id	**G24KA241**	**G24KA311**	**G24KB243**	**G24KB313**
Direction	**Get**	**Get**	**Get**	**Get**
Mode	**ASCII**	**ASCII**	**Binary**	**Binary**
TCP/IP Release	**2.4**	**310**	**2.4**	**310**
File Size (Kb)	24.3	24.3	24.3	24.3
Files Transferred	100	100	100	100
Total Kb	2430	2430	2430	2430
Elapsed Time (sec)	19.2	19.5	15.4	14.7
Rate (Kb/sec)	126.6	124.6	157.8	165.3
Rate Ratio	1.000	0.985	1.000	1.048
TCPIP CPU Usage				
CP CPU (sec)	0.95	0.96	0.94	0.94
Virtual CPU (sec)	1.44	1.38	1.41	1.34
Total CPU (sec)	2.39	2.34	2.35	2.28
Total CPU Ratio	1.000	0.979	1.000	0.970
Total CPU/Kb (msec)	0.984	0.963	0.967	0.938
VM Client CPU Usage				
CP CPU (sec)	0.92	0.90	0.87	0.88
Virtual CPU (sec)	2.93	2.84	1.97	1.88
Total CPU (sec)	3.85	3.74	2.84	2.76
Total CPU Ratio	1.000	0.971	1.000	0.972
Total CPU/Kb (msec)	1.584	1.539	1.169	1.136
Total CPU Usage				
Total CPU (sec)	6.24	6.08	5.19	5.04
Total CPU Ratio	1.000	0.974	1.000	0.971
Total CPU/Kb (msec)	2.568	2.502	2.136	2.074

Table 12. FTP Performance on a 9121-480: Put 2MB and 24KB Binary Files				
Run Id	**P2M_B243**	**P2M_B313**	**P24KB243**	**P24KB313**
Direction	**Put**	**Put**	**Put**	**Put**
Mode	**Binary**	**Binary**	**Binary**	**Binary**
TCP/IP Release	**2.4**	**310**	**2.4**	**310**
File Size (Kb)	2073.8	2073.8	24.3	24.3
Files Transferred	5	5	100	100
Total Kb	10369	10369	2430	2430
Elapsed Time (sec)	29.6	29.6	8.1	8.2
Rate (Kb/sec)	350.3	350.3	300.0	296.3
Rate Ratio	1.000	1.000	1.000	0.988
TCPIP CPU Usage				
CP CPU (sec)	2.06	2.10	1.14	1.13
Virtual CPU (sec)	3.12	2.95	1.67	1.54
Total CPU (sec)	5.18	5.05	2.81	2.67
Total CPU Ratio	1.000	0.975	1.000	0.950
Total CPU/Kb (msec)	0.500	0.487	1.156	1.099
VM Client CPU Usage				
CP CPU (sec)	0.48	0.48	0.77	0.78
Virtual CPU (sec)	0.85	0.84	1.90	1.79
Total CPU (sec)	1.33	1.32	2.67	2.57
Total CPU Ratio	1.000	0.992	1.000	0.963
Total CPU/Kb (msec)	0.128	0.127	1.099	1.058
Total CPU Usage				
Total CPU (sec)	6.51	6.37	5.48	5.24
Total CPU Ratio	1.000	0.978	1.000	0.956
Total CPU/Kb (msec)	0.628	0.614	2.255	2.156

The CPU usage data were obtained from the CP QUERY TIME command. The elapsed times were obtained by summing the elapsed times reported for each file transfer by FTP′s console messages.

Measurement Discussion: There was no significant change in file transfer rate. On average, total CPU usage in the TCPIP virtual machine decreased by 3.7%. This is the same degree of improvement that was observed for the Telnet measurements described in “Telnet” on page 33 and arises from the pathlength improvements that were made to the TCP/IP stack in this release.

Total CPU usage in the client virtual machine decreased by 47% for the GET 2MB ASCII case. The specific reason for this large improvement has not been identified. Client total CPU usage decreased by an average of 2.7% for the remaining cases.

Migration from Other VM Releases

The performance results provided in this report apply to migration from VM/ESA 2.2.0. 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*

230 *VM/ESA Version 2 Release 3.0 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 2.1.0 to VM/ESA 2.3.0 will tend to affect VSE guest performance, look at the VM/ESA 2.1.0 to VM/ESA 2.2.0 comparison measurements and the VM/ESA 2.2.0 to VM/ESA 2.3.0 comparison measurements. In each case, use the measurements from the system configuration that best approximates your VM system.

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 41.

Table 13. 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	
VM/ESA 2.2.0	VM/ESA 2.3.0	9121-742	220			
		9121-480	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 can serve as a good source of additional performance information.

Table 14 summarizes the approximate ITR relationships for the CMS-intensive environment for migrations to VM/ESA 2.3.0.

Table 14. Approximate VM/ESA 2.3.0 relative capacity: CMS-intensive environment

Source	Case	ITRR	Notes
VM/SP 5	9221-120	0.91	1,2,5-7
VM/SP 6	9221-120	1.06	2,5-7
VM/ESA 1.0 (370)	9221-120 9221-170	0.99 1.06	2,5-7 4-7
VM/ESA 1.5 (370)	9221-120 9221-170	0.97 1.04	2,5-7 4-7
VM/SP HPO 5	UP MP	1.00 1.11	4,6,7 3,4,6,7
VM/XA 2.0		1.23	7
VM/XA 2.1		1.20	7
VM/ESA 1.0 ESA		1.16	7
VM/ESA 1.1		1.11	7
VM/ESA 2		1.10	7
VM/ESA 2.1		1.09	7
VM/ESA 2.2		1.06	7
VM/ESA 2.1.0		1.01	
VM/ESA 2.2.0		1.00	

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.

ITRR VM/ESA 2.3.0 ITR divided by the source ITR. A number greater than 1.00 indicates an improvement in processor capacity.

Notes Applicable notes (described below).

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. VM/ESA 2.3.0 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 VM/ESA 2.3.0 system has additional real storage.

   For example, the stated VM/SP HPO 5 to VM/ESA 2.3.0 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 14 on page 41 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 2.3.0. From Table 14 on page 41, the estimated ITRR for migrating from VM/ESA 2.1 to VM/ESA 2.3.0 is 1.09. Therefore, the estimated overall increase in system capacity is 1.30*1.09 = 1.42.

Table 14 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.

New Functions

A number of the functional enhancements in VM/ESA 2.3.0 have performance implications. This section contains performance evaluation results for the following functions:

- Record Level Minidisk Caching
- fork()
- MQSeries Client Support

Record Level Minidisk Caching

This section provides measurement results that compare the performance of record level minidisk caching to the (default) track-based minidisk caching. Results are shown for a random access workload, a sequential access workload, and the FS8F0R CMS-intensive workload.

See "Performance Improvements" on page 3 for a description of record level minidisk caching.

Random / Sequential Results

A CMS program was written to do I/Os to a very large minidisk file (about 900 3390 cylinders). The program could be run in either random or sequential mode. When run in random mode, it simulated the mostly-random I/O pattern of a customer database environment. When run in sequential mode, it read the file sequentially.

These two workloads were run on a dedicated 9121-480 configured with 128MB of central storage and 128MB of expanded storage. The real storage minidisk cache and the expanded storage minidisk cache were both fixed at 2MB. The elapsed time results are summarized in Table 15.

Table 15. Record level MDC performance with random and sequential workloads

File Access	Track	Record	Ratio
Random	1353	883	0.65
Sequential	294	587	2.00

Note: Elapsed time in seconds.

These results illustrate that record level caching can result in substantial performance improvements for workloads characterized by random access of large files. However, more sequential workloads or those with a greater locality of reference will tend to do worse with record level minidisk cache.

CMS-Intensive Results

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*
3390-3	RAMAC 2	4		8		16		
3390-2	3990-2	4	16		6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4			2			

Note: *R* next to the DASD counts means basic cache enabled; *W* means DASD fast write (and basic cache) enabled. RAMAC 2 refers to the RAMAC 2 Array Subsystem with 256MB cache and drawers in 3390-3 format. With the RAMAC 2 Array Subsystem, read and write caching are always enabled.

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	2040	Users	3MB/XC	100		

Measurement Discussion: Table 16 on page 46 shows a comparison of record level minidisk cache to full track minidisk cache for the FS8F0R CMS-intensive workload run on the 9121-480 processor. These results show that overall performance was similar for this workload. With record level caching, external response time (AVG LAST(T)) improved by 10% while processor capacity (ITR(H)) decreased by 0.8%.

DASD I/Os (NONPAGE RIO/CMD(V)) increased and MDC HIT RATIO decreased due to the frequent appearance of multiple cache misses per track. However, average DASD access time (DASD RESP TIME(V)) and minidisk cache size (MDC

REAL SIZE (MB)) decreased because unneeded DASD records were not being read into the cache.

The decrease in processor capacity was primarily due to a 1.5% increase in CP CPU usage (CP/CMD(H)), reflecting the fact that record level MDC is somewhat less efficient than full track caching for typical CMS workloads.

The FS8F0R workload generates an I/O access pattern that is intermediate between the random and sequential workloads described earlier.

Table 16 (Page 1 of 3). Record level minidisk cache: FS8F0R CMS-intensive workload on the 9121-480

Minidisk Cache **Release** **Run ID**	**Track** **2.3.0** **L2AE2043**	**Record** **2.3.0** **L2AE2044**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2040**	**2040**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.131	0.135	0.004	3.05%
NONTRIV INT	0.468	0.451	-0.017	-3.63%
TOT INT	0.350	0.343	-0.007	-2.00%
TOT INT ADJ	0.320	0.307	-0.013	-4.06%
AVG FIRST (T)	0.298	0.265	-0.032	-10.76%
AVG LAST (T)	0.424	0.382	-0.042	-9.80%
Throughput				
AVG THINK (T)	26.15	26.17	0.02	0.10%
ETR	65.19	63.90	-1.29	-1.98%
ETR (T)	71.31	71.40	0.09	0.12%
ETR RATIO	0.914	0.895	-0.019	-2.10%
ITR (H)	78.95	78.36	-0.59	-0.75%
ITR	36.09	35.09	-1.01	-2.80%
EMUL ITR	53.91	52.61	-1.30	-2.41%
ITRR (H)	1.000	0.992	-0.008	-0.75%
ITRR	1.000	0.972	-0.028	-2.80%
Proc. Usage				
PBT/CMD (H)	25.331	25.523	0.192	0.76%
PBT/CMD	25.381	25.490	0.109	0.43%
CP/CMD (H)	8.964	9.094	0.130	1.45%
CP/CMD	8.413	8.403	-0.010	-0.12%
EMUL/CMD (H)	16.367	16.429	0.062	0.38%
EMUL/CMD	16.967	17.086	0.119	0.70%
Processor Util.				
TOTAL (H)	180.65	182.24	1.59	0.88%
TOTAL	181.00	182.00	1.00	0.55%
UTIL/PROC (H)	90.33	91.12	0.80	0.88%
UTIL/PROC	90.50	91.00	0.50	0.55%
TOTAL EMUL (H)	116.72	117.31	0.59	0.50%
TOTAL EMUL	121.00	122.00	1.00	0.83%
MASTER TOTAL (H)	90.04	90.85	0.81	0.90%
MASTER TOTAL	90.00	91.00	1.00	1.11%
MASTER EMUL (H)	51.57	51.96	0.39	0.76%
MASTER EMUL	54.00	54.00	0.00	0.00%
TVR(H)	1.55	1.55	0.01	0.38%
TVR	1.50	1.49	0.00	-0.27%

Table 16 (Page 2 of 3). Record level minidisk cache: FS8F0R CMS-intensive workload on the 9121-480

Minidisk Cache Release Run ID	**Track 2.3.0 L2AE2043**	**Record 2.3.0 L2AE2044**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2040**	**2040**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Storage				
NUCLEUS SIZE (V)	2452KB	2452KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	83	83	0	0.00%
PGBLPGS	54125	54148	23	0.04%
PGBLPGS/USER	26.5	26.5	0.0	0.04%
TOT PAGES/USER (V)	173	176	3	1.73%
FREEPGS	6252	6246	-6	-0.10%
FREE UTIL	0.94	0.94	0.00	0.10%
SHRPGS	1490	1460	-30	-2.01%
Paging				
READS/SEC	716	717	1	0.14%
WRITES/SEC	474	471	-3	-0.63%
PAGE/CMD	16.687	16.638	-0.048	-0.29%
PAGE IO RATE (V)	200.500	198.700	-1.800	-0.90%
PAGE IO/CMD (V)	2.811	2.783	-0.029	-1.02%
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	8.862	8.865	0.003	0.04%
Queues				
DISPATCH LIST	46.59	44.20	-2.39	-5.13%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	703	710	7	1.00%
VIO/CMD	9.858	9.944	0.086	0.87%
RIO RATE (V)	394	443	49	12.44%
RIO/CMD (V)	5.525	6.204	0.680	12.30%
NONPAGE RIO/CMD (V)	2.713	3.421	0.708	26.10%
DASD RESP TIME (V)	18.700	16.000	-2.700	-14.44%
MDC REAL SIZE (MB)	40.7	17.1	-23.6	-58.01%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	207	204	-3	-1.45%
MDC WRITES (I/Os)	9.70	9.33	-0.37	-3.81%
MDC AVOID	195	173	-22	-11.28%
MDC HIT RATIO	0.94	0.83	-0.11	-11.70%
PRIVOPs				
PRIVOP/CMD	13.867	13.928	0.061	0.44%
DIAG/CMD	26.058	26.132	0.074	0.28%
DIAG 04/CMD	2.258	2.256	-0.003	-0.13%
DIAG 08/CMD	0.735	0.735	0.000	-0.05%
DIAG 0C/CMD	0.192	0.193	0.001	0.44%
DIAG 14/CMD	0.025	0.024	0.000	-0.38%
DIAG 58/CMD	1.249	1.250	0.001	0.11%
DIAG 98/CMD	1.016	1.081	0.065	6.38%
DIAG A4/CMD	3.607	3.610	0.003	0.07%
DIAG A8/CMD	2.659	2.682	0.023	0.87%
DIAG 214/CMD	11.997	11.982	-0.015	-0.12%
DIAG 270/CMD	0.941	0.941	0.001	0.07%
SIE/CMD	53.622	54.438	0.817	1.52%
SIE INTCPT/CMD	34.318	34.296	-0.022	-0.06%
FREE TOTL/CMD	49.485	51.567	2.082	4.21%

Table 16 (Page 3 of 3). Record level minidisk cache: FS8F0R CMS-intensive workload on the 9121-480				
Minidisk Cache **Release** **Run ID**	**Track** **2.3.0** **L2AE2043**	**Record** **2.3.0** **L2AE2044**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2040**	**2040**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
VTAM Machines				
WKSET (V)	567	558	-9	-1.59%
TOT CPU/CMD (V)	3.9496	3.9915	0.0419	1.06%
CP CPU/CMD (V)	1.4801	1.5017	0.0216	1.46%
VIRT CPU/CMD (V)	2.4695	2.4898	0.0203	0.82%
DIAG 98/CMD (V)	1.017	1.081	0.065	6.35%
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

fork()

This section contains measurement results that evaluate the performance of the fork() support that has been added to OpenEdition for VM/ESA in VM/ESA 2.3.0.

The VM/ESA support is a partial implementation of the fork() function. Most notably, the VM implementation does not create a separate address space for the child process. Because of this and other functional differences, fork() on VM/ESA comes with certain restrictions. For example, fork() must be followed by one of the exec functions (execl(), execle(), etc) in the child process. For further information on fork() and the VM/ESA restrictions, see *C for VM/ESA Library Reference*.

A C program was written that executes a loop consisting of fork(), execl() by the child process to another C program that does an immediate return, and wait() by the parent for child process termination. This program was measured with one such loop iteration and with 101 loop iterations. The 1-iteration results were subtracted from the 101-iteration results and then divided by 100 in order to obtain per-iteration performance figures. Elapsed time and CPU usage data for both the user virtual machine and the Byte File System (BFS) server virtual machine were collected using the CP QUERY TIME command. BFS server statistics were collected using the QUERY FILEPOOL COUNTER command. All this information was collected immediately prior to and immediately following execution of the C program.

The fraction of all user virtual machine CPU usage due to fork() and wait() was estimated by analysis of instruction trace data. The remaining virtual machine CPU usage and all of the BFS server activity resulted from execl(), including the loading and execution of the null C program.

The measurements were made on a non-dedicated 2003-156 during low usage conditions when contention from other system activity was minimal. The BFS server was dedicated during the measurement. Multiple measurement runs were obtained to verify repeatability. The results, shown in Table 17, are from a typical measurement run.

Table 17. Performance of individual OpenEdition for VM/ESA calls on a 9021-900 processor						
Function	**Elapsed Time (msec)**	**Total CPU Time (msec)**	**User CPU TIme (msec)**	**Server CPU Time (msec)**	**BFS Calls**	**BFS I/Os**
fork() + wait()	2.3	2.3	2.3	0	0	0
execl()	123.0	35.8	30.9	4.9	3	3

The CPU results have been renormalized to a 9021-900 so that they can be directly compared to the results obtained for other OpenEdition for VM/ESA functions published in the *VM/ESA Version 2 Release 1.0 Performance Report*.

As shown in the table, most of the measured CPU usage is from execl(). The low fork() CPU usage reflects the fact that VM/ESA′s fork() implementation does not create the child process in a new address space.

MQSeries Client Support

The Message Queueing Interface is a popular means of enabling applications on different computing systems/architectures to work together. This interface is now available to VM applications through use of the MQSeries* client support provided by VM/ESA 2.3.0. This support, linked with the VM application, enables messages to be added to and removed from queues managed by an MQSeries queue manager residing on another platform elsewhere in the network. This section summarizes the results of performance measurements that were obtained to evaluate the performance of this support.

Measurement Method: For VM/ESA, the MQ* API is available to applications written in C, COBOL, PL/I, assembler, and REXX. Measurements were obtained for the C and REXX cases. The measured workloads were based on the sample applications provided on the 193 CMS minidisk. These write a given number of messages to a specified queue using an MQPUT loop and then read them all back using an MQGET loop. For the C measurements, AMQSPUT0 was modified to read the messages to be put on the queue from a CMS file. For the REXX measurements, the RXMQVPUT and RXMQVGET samples were modified to comment out all console messages that are produced during execution of the MQGET and MQPUT loops. The RXMQV module was NUCXLOADed before doing the REXX measurements, as recommended by VM documentation.

The client workload was run on a VM/ESA 2.3.0 production system on a 2003-156 processor. The measurements were obtained during a period of low system usage so as to minimize the effects of contention from other concurrent activity. The target queue was managed by an MQ server that was running on an AIX RS/6000 system in a remote location. TCP/IP was used for communications.

Measurements were made for the case of 1 message and for 500 messages. Separate data were collected for the put and get operations. Elapsed time and CPU usage by the virtual machine running the tests were obtained using the CP QUERY TIME command. The 1 message results were subtracted from the 500 message results to factor out initialization and termination and the differences were divided by 499 to express the results on a per-message basis. The results for a typical set of measurements are summarized in Table 18. Repeat measurements gave similar results.

Measurement Results

Table 18. VM MQSeries performance on a 2003-156

Language	Function	Elapsed Time per msg (ms)	CPU Time per msg (ms)
C	get put	65 67	1.0 0.9
REXX	get put	99 103	13.0 14.1

The CPU usage shown in Table 18 is for the virtual machine running the MQ application. It includes CPU usage within the virtual machine itself and CPU usage for CP services used by that virtual machine (TOTCPU). There is additional CPU usage in the TCPIP virtual machine resulting from TCP/IP communications activity. INDICATE USER measurements of the TCPIP virtual machine indicate that there is approximately 2 msec of TCPIP CPU usage per MQ message.

The COBOL, PL/I, and assembler cases, not measured, should have performance characteristics that are similar to the C results shown here. The REXX case has substantially higher CPU requirements. This is due to the more complex interface that is needed to communicate between the REXX program and the RXMQV module that contains the actual MQ client code.

Elapsed time per message is heavily determined by the performance characteristics of the network connecting the client and server systems. Because of this, elapsed times for other network configurations can vary substantially from these results.

For more information on the sample applications, see Appendix C of the "CMS Application Development Guide." For more information about MQSeries, see the "MQSeries: Application Programming Guide" and the other MQSeries publications.

Additional Evaluations

This portion of the report includes results from a number of additional VM/ESA (including TCP/IP) performance measurement evaluations that have been conducted during the past year.

Migration from VTAM to Telnet

The section explores the performance implications of migrating end-user 3270 connectivity from VTAM to TCP/IP Telnet.

Measurements were obtained by running the FS8F0R workload on a 9121-480 processor with the end users simulated by TPNS running on a separate system. VM/ESA 2.2.0 was used for all measurements. For the base measurement, connectivity was provided by VTAM 3.4.1 through a CTCA connection with the TPNS system. This measurement configuration is equivalent to that used for the CMS-intensive measurements shown in “9121-480 / Minidisk” on page 17.

Table 19 compares this VTAM base measurement to a corresponding measurement using TCP/IP 2.4 Telnet through a CTCA connection, while Table 20 compares the VTAM base measurement to a corresponding TCP/IP 2.4 measurement with connectivity to the TPNS system provided through a 3172-3 Interconnect Controller and a 16Mbit IBM Token Ring. For each measurement, the number of simulated users was chosen such that the measured 9121-480 system was running at approximately 90% processor utilization.

Workload: FS8F0R

Hardware Configuration

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

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-3	RAMAC 2	4		8		16		
3390-2	3990-2	4	16		6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4			2			

Note: *R* next to the DASD counts means basic cache enabled; *W* means DASD fast write (and basic cache) enabled. RAMAC 2 refers to the RAMAC 2 Array Subsystem with 256MB cache and drawers in 3390-3 format. With the RAMAC 2 Array Subsystem, read and write caching are always enabled.

Communications (CTCA):

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

Communications (Token Ring):

16 Mbit IBM Token Ring
3172-3 Interconnect Controller

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*
VTAMXA or	1	VTAM/VSCS	64MB/XA	10000	560	QUICKDSP ON
TCPIP	1	TCP/IP	256MB/XA	10000	2700	QUICKDSP ON
SMART	1	RTM	32MB/XA	3%	400	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	2000/1800	Users	3MB/XC	100		

Measurement Discussion: The Table 19 comparison is more equivalent because the method of physical interconnection is held constant. The Table 20 comparison shows TCP/IP Telnet performance for the more typical LAN-based interconnection.

Looking at the Table 19 comparison, we see that external response time (AVG LAST(T)) decreased by 4%, while the internal throughput rate (ITR(H)) decreased by 10%. The ITR drop reflects the fact that TCP/IP CPU usage is somewhat higher than VTAM′s for this type of function. Roughly half of this increase shows up as processing time that is charged to the TCPIP virtual machine, which includes both the TCP/IP protocol stack and internal Telnet client code. The remaining increase shows up in CP CPU usage that is charged to the end users as the result of Diagnose Code X′7C′ processing, which operates on behalf of both the TCPIP virtual machine and the user virtual machines for which it is providing logical device support.

VTAM supports the 3270 interface through the *CCS CP system service (accessed using IUCV requests), while Telnet provides this function through use of the Diagnose X′7C′ logical device support facility. This difference is reflected in the results as a large decrease in PRIVOP/CMD and much of the large increase in DIAG/CMD. The fact that diagnose X′7C′ has a longer pathlength than *CCS accounts for much of the CPU usage increase observed in the TCP/IP measurement relative to the VTAM base measurement. Another contributing factor is that TCP/IP does more communication I/Os than VTAM, as shown by the increase in DIAG 98/CMD(V).

The TCPIP virtual machine uses more real storage than does the VTAM virtual machine (see WKSET(V)). More pages had to be reserved in the TCPIP case (2700 as opposed to 560 for VTAM) in order to eliminate paging from the communications server. The “trimmed” buffer pool configuration defined in Table 21 on page 59 was used for the TCP/IP measurements. As discussed in “Reduced TCP/IP Real Storage Requirements” on page 58, it was necessary for these TCP/IP 2.4 measurements to avoid defining an excessive number of TCP/IP buffers but this tuning is no longer necessary when using TCP/IP 310.

For the most part, the above discussion also applies to the LAN-based Table 20 comparison. The main differences are that TCP/IP response times are better and the amount of CPU usage increase relative to the VTAM base measurement is somewhat smaller when the VTAM base measurement is compared to TCP/IP Telnet with token ring interconnection.

Until recently, the number of Telnet users on a given VM/ESA system was limited to 4096 — the maximum number of logical devices supported by CP. With VM/ESA 2.3.0, this limit has been raised to 32,768.

The observations shown here should also apply to larger systems with larger numbers of users. However, CPU usage in the TCPIP virtual machine can become a bottleneck on large systems. If we start with the 27.9% TCPIP CPU utilization observed for run L29E1801, a simple linear projection indicates that a system with a per-engine processor capacity similar to the 9121-480 would bottleneck on TCPIP CPU usage at or below 1800*(100/28) = 6400 users. Much of VTAM′s processor usage can be offloaded to one or more VSCS virtual machines but TCP/IP VM does not have a corresponding capability.

Table 19 (Page 1 of 2). VTAM to Telnet Migration on a 9121-480: CTCA

Communications Interconnection Run ID	**VTAM 3.4.1 CTCA L29E2006**	**TCP/IP 2.4 CTCA L29E1800**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2000**	**1800**	**-200**	**-10.00%**
VTAMs	**1**	**na**		
VSCSs	**0**	**na**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.127	0.114	-0.013	-10.24%
NONTRIV INT	0.441	0.861	0.420	95.24%
TOT INT	0.331	0.246	-0.085	-25.68%
TOT INT ADJ	0.301	0.225	-0.076	-25.14%
AVG FIRST (T)	0.274	0.293	0.020	7.13%
AVG LAST (T)	0.390	0.404	0.015	3.72%
Throughput				
AVG THINK (T)	26.17	24.51	-1.66	-6.33%
ETR	63.65	57.72	-5.93	-9.32%
ETR (T)	70.00	63.02	-6.98	-9.97%
ETR RATIO	0.909	0.916	0.007	0.73%
ITR (H)	79.41	71.11	-8.30	-10.46%
ITR	36.14	32.59	-3.55	-9.82%
EMUL ITR	53.89	53.27	-0.62	-1.15%
ITRR (H)	1.000	0.895	-0.105	-10.46%
ITRR	1.000	0.902	-0.098	-9.82%
Proc. Usage				
PBT/CMD (H)	25.186	28.127	2.941	11.68%
PBT/CMD	25.143	28.086	2.944	11.71%
CP/CMD (H)	8.891	11.579	2.688	30.24%
CP/CMD	8.286	10.949	2.663	32.14%
EMUL/CMD (H)	16.295	16.548	0.253	1.55%
EMUL/CMD	16.857	17.138	0.280	1.66%
Processor Util.				
TOTAL (H)	176.30	177.25	0.95	0.54%
TOTAL	176.00	177.00	1.00	0.57%
UTIL/PROC (H)	88.15	88.63	0.48	0.54%
UTIL/PROC	88.00	88.50	0.50	0.57%
TOTAL EMUL (H)	114.06	104.28	-9.78	-8.58%
TOTAL EMUL	118.00	108.00	-10.00	-8.47%
MASTER TOTAL (H)	87.68	89.26	1.58	1.81%
MASTER TOTAL	88.00	89.00	1.00	1.14%
MASTER EMUL (H)	50.08	43.08	-7.00	-13.98%
MASTER EMUL	52.00	45.00	-7.00	-13.46%
TVR(H)	1.55	1.70	0.15	9.97%
TVR	1.49	1.64	0.15	9.88%
Storage				
NUCLEUS SIZE (V)	2804KB	2804KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	82	84	2	2.44%
PGBLPGS	54314	55503	1189	2.19%
PGBLPGS/USER	27.2	30.8	3.7	13.54%
TOT PAGES/USER (V)	172	177	5	2.91%
FREEPGS	5964	5206	-758	-12.71%
FREE UTIL	0.94	0.98	0.04	4.15%
SHRPGS	1465	946	-519	-35.43%

Table 19 (Page 2 of 2). VTAM to Telnet Migration on a 9121-480: CTCA

Communications Interconnection Run ID	**VTAM 3.4.1 CTCA L29E2006**	**TCP/IP 2.4 CTCA L29E1800**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2000**	**1800**	**-200**	**-10.00%**
VTAMs	**1**	**na**		
VSCSs	**0**	**na**		
Processors	**2**	**2**		
Paging				
READS/SEC	697	660	-37	-5.31%
WRITES/SEC	458	408	-50	-10.92%
PAGE/CMD	16.500	16.947	0.447	2.71%
PAGE IO RATE (V)	191.400	170.300	-21.100	-11.02%
PAGE IO/CMD (V)	2.734	2.702	-0.032	-1.17%
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	8.757	8.759	0.002	0.02%
Queues				
DISPATCH LIST	46.04	23.49	-22.55	-48.97%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	692	875	183	26.45%
VIO/CMD	9.886	13.885	3.999	40.45%
RIO RATE (V)	388	608	220	56.70%
RIO/CMD (V)	5.543	9.648	4.105	74.06%
NONPAGE RIO/CMD (V)	2.809	6.945	4.137	147.29%
DASD RESP TIME (V)	19.800	19.700	-0.100	-0.51%
MDC REAL SIZE (MB)	41.0	42.6	1.6	3.88%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	199	179	-20	-10.05%
MDC WRITES (I/Os)	9.51	8.69	-0.82	-8.62%
MDC AVOID	188	168	-20	-10.64%
MDC HIT RATIO	0.94	0.94	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.942	1.603	-12.339	-88.50%
DIAG/CMD	26.668	36.640	9.972	37.39%
DIAG 04/CMD	2.373	2.582	0.209	8.80%
DIAG 08/CMD	0.735	0.733	-0.003	-0.37%
DIAG 0C/CMD	0.212	0.212	0.000	0.02%
DIAG 14/CMD	0.025	0.024	0.000	-1.24%
DIAG 58/CMD	1.251	1.251	0.001	0.06%
DIAG 7C/CMD	0.000	5.783	5.783	na
DIAG 98/CMD	1.082	5.102	4.020	371.47%
DIAG A4/CMD	3.560	3.546	-0.013	-0.38%
DIAG A8/CMD	2.672	2.662	-0.011	-0.40%
DIAG 214/CMD	12.460	12.444	-0.017	-0.13%
DIAG 270/CMD	0.920	0.918	-0.002	-0.25%
SIE/CMD	52.929	64.424	11.496	21.72%
SIE INTCPT/CMD	34.404	44.453	10.049	29.21%
FREE TOTL/CMD	49.629	53.682	4.053	8.17%
VTAM or TCPIP Machine				
WKSET (V)	547	2700	2153	393.60%
CPU UTIL (V)	27.7	33.1	5.4	19.49%
TOT CPU/CMD (V)	3.9524	5.2360	1.2836	32.48%
CP CPU/CMD (V)	1.4762	2.3640	0.8878	60.14%
VIRT CPU/CMD (V)	2.4762	2.8720	0.3958	15.98%
DIAG 98/CMD (V)	1.082	5.102	4.020	371.62%

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

Table 20 (Page 1 of 2). VTAM/CTCA to Telnet/Token Ring Migration on a 9121-480

Communications Interconnection Run ID	**VTAM 3.4.1 CTCA L29E2006**	**TCP/IP 2.4 3172-3/TR L29E1801**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2000**	**1800**	**-200**	**-10.00%**
VTAMs	**1**	**na**		
VSCSs	**0**	**na**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.127	0.093	-0.034	-26.77%
NONTRIV INT	0.441	0.804	0.363	82.31%
TOT INT	0.331	0.221	-0.110	-33.23%
TOT INT ADJ	0.301	0.198	-0.103	-34.15%
AVG FIRST (T)	0.274	0.205	-0.068	-24.86%
AVG LAST (T)	0.390	0.306	-0.084	-21.44%
Throughput				
AVG THINK (T)	26.17	24.48	-1.68	-6.42%
ETR	63.65	56.59	-7.06	-11.09%
ETR (T)	70.00	63.10	-6.90	-9.86%
ETR RATIO	0.909	0.897	-0.012	-1.37%
ITR (H)	79.41	73.40	-6.01	-7.56%
ITR	36.14	32.94	-3.19	-8.83%
EMUL ITR	53.89	53.50	-0.39	-0.73%
ITRR (H)	1.000	0.924	-0.076	-7.56%
ITRR	1.000	0.912	-0.088	-8.83%
Proc. Usage				
PBT/CMD (H)	25.186	27.246	2.061	8.18%
PBT/CMD	25.143	27.258	2.115	8.41%
CP/CMD (H)	8.891	11.083	2.192	24.66%
CP/CMD	8.286	10.459	2.174	26.23%
EMUL/CMD (H)	16.295	16.163	-0.131	-0.81%
EMUL/CMD	16.857	16.799	-0.059	-0.35%
Processor Util.				
TOTAL (H)	176.30	171.93	-4.37	-2.48%
TOTAL	176.00	172.00	-4.00	-2.27%
UTIL/PROC (H)	88.15	85.96	-2.19	-2.48%
UTIL/PROC	88.00	86.00	-2.00	-2.27%
TOTAL EMUL (H)	114.06	101.99	-12.07	-10.58%
TOTAL EMUL	118.00	106.00	-12.00	-10.17%
MASTER TOTAL (H)	87.68	86.71	-0.96	-1.10%
MASTER TOTAL	88.00	87.00	-1.00	-1.14%
MASTER EMUL (H)	50.08	41.61	-8.47	-16.92%
MASTER EMUL	52.00	43.00	-9.00	-17.31%
TVR(H)	1.55	1.69	0.14	9.06%
TVR	1.49	1.62	0.13	8.79%
Storage				
NUCLEUS SIZE (V)	2804KB	2804KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	82	86	4	4.88%
PGBLPGS	54314	55481	1167	2.15%
PGBLPGS/USER	27.2	30.8	3.7	13.50%
TOT PAGES/USER (V)	172	177	5	2.91%
FREEPGS	5964	5229	-735	-12.32%
FREE UTIL	0.94	0.98	0.03	3.69%
SHRPGS	1465	897	-568	-38.77%

Table 20 (Page 2 of 2). VTAM/CTCA to Telnet/Token Ring Migration on a 9121-480

Communications Interconnection Run ID	**VTAM 3.4.1 CTCA L29E2006**	**TCP/IP 2.4 3172-3/TR L29E1801**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**2000**	**1800**	**-200**	**-10.00%**
VTAMs	**1**	**na**		
VSCSs	**0**	**na**		
Processors	**2**	**2**		
Paging				
READS/SEC	697	663	-34	-4.88%
WRITES/SEC	458	408	-50	-10.92%
PAGE/CMD	16.500	16.973	0.473	2.87%
PAGE IO RATE (V)	191.400	169.100	-22.300	-11.65%
PAGE IO/CMD (V)	2.734	2.680	-0.054	-1.99%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	8.757	8.780	0.022	0.26%
Queues				
DISPATCH LIST	46.04	22.07	-23.97	-52.06%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	692	740	48	6.94%
VIO/CMD	9.886	11.727	1.842	18.63%
RIO RATE (V)	388	470	82	21.13%
RIO/CMD (V)	5.543	7.448	1.906	34.38%
NONPAGE RIO/CMD (V)	2.809	4.769	1.960	69.79%
DASD RESP TIME (V)	19.800	19.800	0.000	0.00%
MDC REAL SIZE (MB)	41.0	42.5	1.4	3.52%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	199	180	-19	-9.55%
MDC WRITES (I/Os)	9.51	8.69	-0.82	-8.62%
MDC AVOID	188	170	-18	-9.57%
MDC HIT RATIO	0.94	0.94	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.942	1.586	-12.356	-88.62%
DIAG/CMD	26.668	34.642	7.973	29.90%
DIAG 04/CMD	2.373	2.576	0.203	8.54%
DIAG 08/CMD	0.735	0.737	0.001	0.20%
DIAG 0C/CMD	0.212	0.212	0.001	0.33%
DIAG 14/CMD	0.025	0.025	0.000	0.03%
DIAG 58/CMD	1.251	1.250	-0.001	-0.05%
DIAG 7C/CMD	0.000	5.880	5.880	na
DIAG 98/CMD	1.082	2.916	1.834	169.46%
DIAG A4/CMD	3.560	3.566	0.007	0.18%
DIAG A8/CMD	2.672	2.680	0.008	0.29%
DIAG 214/CMD	12.460	12.502	0.042	0.33%
DIAG 270/CMD	0.920	0.920	-0.001	-0.08%
SIE/CMD	52.929	56.988	4.060	7.67%
SIE INTCPT/CMD	34.404	38.752	4.348	12.64%
FREE TOTL/CMD	49.629	51.568	1.940	3.91%
VTAM or TCPIP Machine				
WKSET (V)	547	2700	2153	393.60%
CPU UTIL (V)	27.7	27.9	0.2	0.72%
TOT CPU/CMD (V)	3.9524	4.4290	0.4766	12.06%
CP CPU/CMD (V)	1.4762	1.9990	0.5228	35.42%
VIRT CPU/CMD (V)	2.4762	2.4300	-0.0462	-1.87%
DIAG 98/CMD (V)	1.082	2.916	1.834	169.55%

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

Reduced TCP/IP Real Storage Requirements

The page reference locality within the TCPIP virtual machine has been significantly improved with TCP/IP 310, especially for Telnet. In prior releases, excess buffers were periodically referenced even though they were never actually used. This tended to result in page faults in the TCPIP virtual machine, degrading performance. TCP/IP 310 fixes this problem. In addition, certain buffers and control blocks have been reorganized in order to occupy less storage. As a result, performance is no longer adversely affected by the presence of unused TCP/IP buffers. This is shown by the performance results provided in this section.

Workload: FS8F0R

Hardware Configuration

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

DASD:

Type of DASD	*Control Unit*	*Number of Paths*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-3	RAMAC 2	4		8		16		
3390-2	3990-2	4	16		6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4			2			

Note: *R* next to the DASD counts means basic cache enabled; *W* means DASD fast write (and basic cache) enabled. RAMAC 2 refers to the RAMAC 2 Array Subsystem with 256MB cache and drawers in 3390-3 format. With the RAMAC 2 Array Subsystem, read and write caching are always enabled.

Communications:

16 Mbit IBM Token Ring
3172-3 Interconnect Controller

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*
TCPIP	1	TCP/IP	256MB/XA	10000	2700	QUICKDSP ON
SMART	1	RTM	32MB/XA	3%	400	QUICKDSP ON
WRITER	1	CP monitor	2MB/XA	100		QUICKDSP ON
Unnnn	1800	Users	3MB/XC	100		

Two different TCP/IP buffer pool configurations were used (see Table 21). The "trimmed" case was determined by setting the size of each buffer pool to be about 10% larger than the high water mark requirements shown by the NETSTAT POOLSIZE command while running the Telnet workload.

Table 21. Measured buffer pool settings

TCP/IP buffer pool	Excess	Trimmed
ACBPoolSize	10000	1000
AddressTranslationPoolSize	1500	150
CCBPoolSize	500	20
DataBufferPoolSize	15000 8192	150 8192
EnvelopePoolSize	5000	300
IPRoutePoolSize	300	100
LargeEnvelopePoolSize	500	10
RCBPoolSize	10	10
SCBPoolSize	5000	20
SKCBPoolSize	5000	10
SmallDataBufferPoolSize	15000 2048	2500 2048
TCBPoolSize	10000	2000
TinyDataBufferPoolSize	5000	20
UCBPoolSize	10	10

Measurement Discussion: Preliminary TCP/IP 2.4 measurements using the excess buffer pool configuration yielded results that showed severe performance degradation due to high amounts of paging in the TCPIP virtual machine even though 2700 pages were reserved. In one example measurement, average response time was 4.6 seconds and the TCPIP virtual machine was sustaining 52 page reads per second. This problem went away when the trimmed buffer pool configuration was used. The TCPIP virtual machine then ran with zero paging when 2700 pages were reserved for it (see the first measurement in Table 22).

In TCP/IP 310, several aspects of TCPIP buffer pool management were changed in order to solve this page reference problem. In order to verify the effectiveness of these changes, the TCP/IP 2.4 measurement with the trimmed buffer pool configuration was rerun using TCP/IP 310 and the original excess buffer pool configuration (the second measurement in Table 22). The two measurements are equivalent within run variability.

Table 22 (Page 1 of 2). Telnet results showing improved TCP/IP buffer management

TCP/IP Release **TCP/IP Buffers** **Run ID**	**2.4** **Trimmed** **L29E1800**	**310** **Excess** **L29E1805**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1800**	**1800**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.114	0.109	-0.005	-4.39%
NONTRIV INT	0.861	0.859	-0.002	-0.23%
TOT INT	0.246	0.239	-0.007	-2.85%
TOT INT ADJ	0.225	0.220	-0.005	-2.20%
AVG FIRST (T)	0.293	0.296	0.003	1.02%
AVG LAST (T)	0.404	0.404	0.000	0.12%
Throughput				
AVG THINK (T)	24.51	24.49	-0.02	-0.08%
ETR	57.72	57.97	0.25	0.43%
ETR (T)	63.02	62.88	-0.14	-0.23%
ETR RATIO	0.916	0.922	0.006	0.66%
ITR (H)	71.11	70.79	-0.31	-0.44%
ITR	32.59	32.67	0.08	0.25%
EMUL ITR	53.27	53.30	0.03	0.05%
ITRR (H)	1.000	0.996	-0.004	-0.44%
ITRR	1.000	1.003	0.003	0.25%
Proc. Usage				
PBT/CMD (H)	28.127	28.251	0.124	0.44%
PBT/CMD	28.086	28.310	0.224	0.80%
CP/CMD (H)	11.579	11.587	0.008	0.07%
CP/CMD	10.949	10.974	0.025	0.23%
EMUL/CMD (H)	16.548	16.663	0.116	0.70%
EMUL/CMD	17.138	17.336	0.198	1.16%
Processor Util.				
TOTAL (H)	177.25	177.63	0.37	0.21%
TOTAL	177.00	178.00	1.00	0.56%
UTIL/PROC (H)	88.63	88.81	0.19	0.21%
UTIL/PROC	88.50	89.00	0.50	0.56%
TOTAL EMUL (H)	104.28	104.77	0.49	0.47%
TOTAL EMUL	108.00	109.00	1.00	0.93%
MASTER TOTAL (H)	89.26	89.48	0.22	0.24%
MASTER TOTAL	89.00	89.00	0.00	0.00%
MASTER EMUL (H)	43.08	43.42	0.34	0.78%
MASTER EMUL	45.00	45.00	0.00	0.00%
TVR(H)	1.70	1.70	0.00	-0.26%
TVR	1.64	1.63	-0.01	-0.36%
Storage				
NUCLEUS SIZE (V)	2804KB	2804KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	84	85	1	1.19%
PGBLPGS	55503	55513	10	0.02%
PGBLPGS/USER	30.8	30.8	0.0	0.02%
TOT PAGES/USER (V)	177	195	18	10.17%
FREEPGS	5206	5196	-10	-0.19%
FREE UTIL	0.98	0.99	0.00	0.19%
SHRPGS	946	927	-19	-2.01%

Table 22 (Page 2 of 2). Telnet results showing improved TCP/IP buffer management

TCP/IP Release **TCP/IP Buffers** **Run ID**	**2.4** **Trimmed** **L29E1800**	**310** **Excess** **L29E1805**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1800**	**1800**		
Processors	**2**	**2**		
Paging				
READS/SEC	660	655	-5	-0.76%
WRITES/SEC	408	407	-1	-0.25%
PAGE/CMD	16.947	16.891	-0.056	-0.33%
PAGE IO RATE (V)	170.300	167.900	-2.400	-1.41%
PAGE IO/CMD (V)	2.702	2.670	-0.032	-1.18%
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	8.759	8.795	0.036	0.41%
Queues				
DISPATCH LIST	23.49	24.21	0.72	3.06%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	875	875	0	0.00%
VIO/CMD	13.885	13.916	0.032	0.23%
RIO RATE (V)	608	605	-3	-0.49%
RIO/CMD (V)	9.648	9.622	-0.026	-0.26%
NONPAGE RIO/CMD (V)	6.945	6.952	0.006	0.09%
DASD RESP TIME (V)	19.700	19.600	-0.100	-0.51%
MDC REAL SIZE (MB)	42.6	43.0	0.4	0.94%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	179	179	0	0.00%
MDC WRITES (I/Os)	8.69	8.63	-0.06	-0.69%
MDC AVOID	168	169	1	0.60%
MDC HIT RATIO	0.94	0.94	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	1.603	1.611	0.007	0.46%
DIAG/CMD	36.640	36.694	0.054	0.15%
DIAG 04/CMD	2.582	2.585	0.003	0.12%
DIAG 08/CMD	0.733	0.734	0.001	0.16%
DIAG 0C/CMD	0.212	0.212	0.000	-0.02%
DIAG 14/CMD	0.024	0.024	0.000	0.59%
DIAG 58/CMD	1.251	1.250	-0.001	-0.10%
DIAG 98/CMD	5.102	5.117	0.015	0.29%
DIAG A4/CMD	3.546	3.562	0.016	0.44%
DIAG A8/CMD	2.662	2.673	0.011	0.42%
DIAG 214/CMD	12.444	12.459	0.015	0.12%
DIAG 270/CMD	0.918	0.920	0.001	0.16%
SIE/CMD	64.424	64.541	0.116	0.18%
SIE INTCPT/CMD	44.453	43.888	-0.565	-1.27%
FREE TOTL/CMD	53.682	53.630	-0.052	-0.10%
TCPIP Machine				
WKSET (V)	2700	2700	0	0.00%
DASD PAGE RATE (V)	0	0	0	0.00%
TOT CPU/CMD (V)	5.2360	5.3550	0.1190	2.27%
CP CPU/CMD (V)	2.3640	2.3950	0.0310	1.31%
VIRT CPU/CMD (V)	2.8720	2.9600	0.0880	3.06%
DIAG 98/CMD (V)	5.102	5.117	0.015	0.29%

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

Improved SMTP Server Capacity

The maximum throughput that can be handled by any given SMTP server virtual machine has been substantially increased with TCP/IP 310 (see Figure 1). For the workload and system configuration described in this section, throughput increase multiples ranging from 1.6 to 3.4 relative to TCP/IP 2.4 were observed. These improvements resulted from eliminating many of the SMTP 191 minidisk I/Os and from using the *SPL CP system service to read spool files asynchronously.

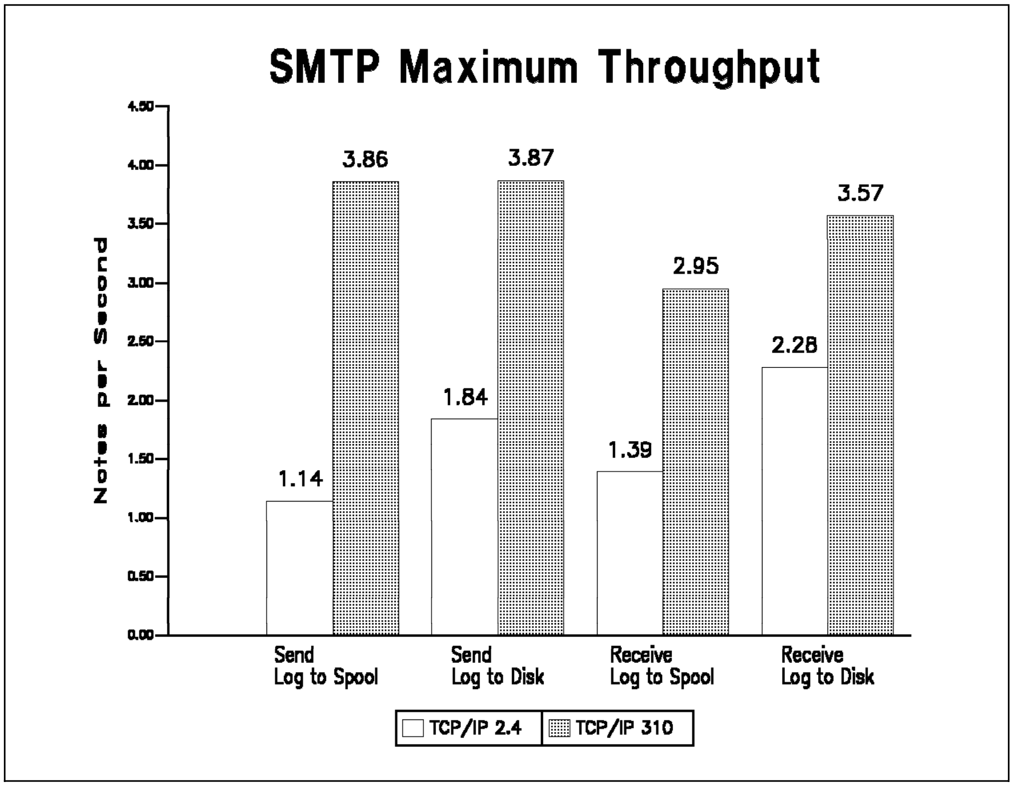

Figure 1. SMTP throughput improvements

Workload: Each note consisted of 60 lines of text and was sent to three recipients. These recipients were all on the same system, which was different from the sending system.

Configuration: The measured system consisted of a 9121-480 (2-way) processor with 256MB of central storage and no expanded storage. It was running VM/ESA 2.3.0 and TCP/IP 2.4 except for SMTP MODULE, which was either at the 2.4 or 310 level. The SMTP 191 minidisk was on a dedicated 3390 volume behind a 3990-3 control unit with read caching but no DASD fast write. VM/ESA minidisk caching was in effect. The 8 spool volumes were on RAMAC 2 Array Subsystem DASD.

The measured system was attached to a 16Mbit IBM Token Ring through a 3172-3 Interconnect Controller. A packet (MTU) size of 4000 bytes was specified. The test notes were sent to or received from one to four additional systems on the same token ring. Two of these additional systems were also VM/ESA 2.3.0 systems with 3172-3 attachment to the token ring. The other two additional systems were RS/6000 model 250s running AIX 4.1.4.

Measurement Method: All note traffic through the measured system was handled by one SMTP server virtual machine. For any given measurement, all notes were either sent from the measured system (send runs) out to recipients on the additional systems or all notes were sent from the additional systems to recipients on the measured system (receive runs).

Multiple additional systems were used (instead of just one) because throughput between any two systems is often limited by network latency. This is because SMTP maintains just one TCP connection between any two systems and communications are serialized on that connection. Of primary interest in these measurements was maximum throughput that can be delivered by one SMTP server when not constrained by the number of connections to other systems.

For send runs, note files in batch SMTP format were punched to the SMTP server from a separate virtual machine on the measured system. Any given note was sent to three recipients on one of the target systems. Notes were sent with equal frequency to all used target systems in order to evenly distribute the load. The number of target systems was varied from run to run in order to determine how many were necessary such that the number of connections did not limit maximum throughput. The applied load was varied from run to run by adjusting the rate at which the sending virtual machine would punch notes to the SMTP server.

For receive runs, note files were sent from one or more of the additional systems to three recipients on the measured system. Notes arrived with roughly equal frequency from each of the sending systems. As with the send runs, the number of sending systems was varied from run to run in order to ensure that throughput was not limited by having too few connections. The applied load was varied from run to run by adjusting the rate at which each of the sending systems would transmit notes to the measured system.

All measurements were taken over a 3 minute interval after the desired level of note traffic had been started. The primary measure was the number of notes delivered per second to the recipients. This was determined by counting the number of “Delivered Note” messages in the SMTP log that occurred during the measurement interval and dividing that by the measurement interval times 3 (the number of recipients per note). That is, one note delivered to its 3 recipients was counted as one note delivered.

Maximum SMTP throughput is significantly affected by whether the SMTP log is to disk or spool (virtual console) because this affects the SMTP minidisk I/Os that are done. Accordingly, maximum throughput was determined for both cases. Except for log placement, the default SMTP CONFIG settings were used.

Measurement Results: The measurement results are summarized in Table 23 and Table 24.

Table 23. SMTP Maximum Throughput Results: Send Notes

runid **SMTP level** **SMTP log** **receiving systems**	**S4SR3022** **2.4** **spool** **3**	**S4SR3021** **2.4** **disk** **3**	**SHSR4053** **310** **spool** **4**	**SHSR4052** **310** **disk** **4**
delivered notes/sec	1.14	1.84	3.86	3.87
ratio vs 2.4 log to spool	1.00	1.61	3.39	3.39
ratio vs 2.4 log to disk	0.62	1.00	2.10	2.10
SMTP CPU-sec/sec	0.11	0.15	0.29	0.27
NAMESRV CPU-sec/sec	0.07	0.11	0.25	0.25
TCPIP CPU-sec/sec	0.06	0.09	0.18	0.18
ms SMTP TCPU/note	96.8	81.1	75.6	71.0
ratio vs 2.4 log to spool	1.00	0.84	0.78	0.73
ratio vs 2.4 log to disk	1.19	1.00	0.93	0.88
ms NAMESRV TCPU/note	60.8	62.3	65.7	65.3
ms TCPIP TCPU/note	52.4	48.6	47.4	47.5
ms TOTAL TCPU/note	212.2	193.8	190.4	185.4
SMTP virtual IOs/note	58.3	43.9	9.7	9.6
spool IOs/sec	14.9	21.2	27.4	26.1
% spool read	54	62	45	47
msec/spool IO	2.5	2.8	3.2	3.0
minidisk IOs/sec	47	51	36	36
msec/minidisk IO	17	14	16	17
minidisk IOs/note	41.2	27.7	9.3	9.3
ratio vs 2.4 log to spool	1.00	0.67	0.23	0.23
ratio vs 2.4 log to disk	1.49	1.00	0.34	0.34
minidisk IO-sec/note	0.70	0.39	0.15	0.16
ratio vs 2.4 log to spool	1.00	0.55	0.21	0.23
ratio vs 2.4 log to disk	1.81	1.00	0.38	0.41
spool read IOs/sec	8.0	13.1	12.3	12.3
spool read IOs/note	7.1	7.1	3.2	3.2
spool read IO-sec/sec	0.02	0.04	0.04	0.04
synch spool read IO-sec/sec	0.02	0.04	0.00	0.00
spool write IOs/sec	6.9	8.1	15.1	13.8
spool write IOs/note	6.0	4.4	3.9	3.6
spool write IO-sec/sec	0.02	0.02	0.05	0.04
synch spool IO-sec/sec	0.04	0.06	0.05	0.04
minidisk IO-sec/sec	0.80	0.71	0.58	0.61
SMTP CPU-sec/sec	0.11	0.15	0.29	0.27
SMTP waiting for CPU-sec/sec	0.02	0.02	0.04	0.04
serialized SMTP-sec/sec	0.97	0.94	0.96	0.96

Table 24. SMTP Maximum Throughput Results: Receive Notes

runid SMTP level SMTP log sending systems	S4RR1041 2.4 spool 1	S4RRD041 2.4 disk 1	SHRR3023 310 spool 3	SHRR3022 310 disk 3
delivered notes/sec	1.39	2.28	2.95	3.57
ratio vs 2.4 log to spool	1.00	1.64	2.12	2.57
ratio vs 2.4 log to disk	0.61	1.00	1.30	1.57
SMTP CPU-sec/sec	0.14	0.22	0.33	0.39
NAMESRV CPU-sec/sec	0.01	0.01	0.01	0.02
TCPIP CPU-sec/sec	0.04	0.06	0.10	0.12
ms SMTP TCPU/note	98.5	96.0	113.0	108.5
ratio vs 2.4 log to spool	1.00	0.97	1.15	1.10
ratio vs 2.4 log to disk	1.03	1.00	1.18	1.13
ms NAMESRV TCPU/note	4.8	5.2	5.0	5.0
ms TCPIP TCPU/note	26.9	26.4	34.8	33.0
ms TOTAL TCPU/note	132.2	129.6	155.1	148.8
SMTP virtual IOs/note	43.4	31.3	8.2	5.1
spool IOs/sec	9.2	14.8	19.3	22.2
% spool read	0.1	0	2	2
msec/spool IO	5.5	5.5	5.5	5.5
minidisk IOs/sec	40	36	23	18
msec/minidisk IO	16	14	17	15
minidisk IOs/note	28.8	15.8	7.8	5.0
ratio vs 2.4 log to spool	1.00	0.55	0.27	0.18
ratio vs 2.4 log to disk	1.82	1.00	0.49	0.32
minidisk IO-sec/note	0.46	0.22	0.13	0.08
ratio vs 2.4 log to spool	1.00	0.48	0.29	0.16
ratio vs 2.4 log to disk	2.08	1.00	0.60	0.34
spool read IOs/sec	0.0	0.0	0.4	0.4
spool read IOs/note	0.0	0.0	0.1	0.1
spool read IO-sec/sec	0.00	0.00	0.00	0.00
synch spool read IO-sec/sec	0.00	0.00	0.00	0.00
spool write IOs/sec	9.2	14.8	18.9	21.8
spool write IOs/note	6.6	6.5	6.4	6.1
spool write IO-sec/sec	0.05	0.08	0.10	0.12
synch spool IO-sec/sec	0.05	0.08	0.10	0.12
minidisk IO-sec/sec	0.64	0.50	0.39	0.27
SMTP CPU-sec/sec	0.14	0.22	0.33	0.39
SMTP waiting for CPU-sec/sec	0.01	0.01	0.01	0.01
serialized SMTP-sec/sec	0.83	0.81	0.83	0.79

The CPU usage data were obtained from the CP QUERY TIME command. The SMTP virtual I/O data were obtained from the CP INDICATE USER command. The minidisk and spool real I/O rates and access times were obtained from RTM/ESA.

Calculations:

Minidisk IO-sec/sec was calculated as minidisk IOs/sec x 0.001 x msec/minidisk IO.

Spool read IOs/sec was calculated as spool IOs/sec x 0.01 x % spool read.

Spool read IO-sec/sec was calculated as spool read IOs/sec x 0.001 x msec/spool IO.

Synch spool read IO-sec/sec was calculated as equal to spool read IO-sec/sec for 2.4 SMTP and zero for 310 SMTP.

Spool write calculations were done in the same manner as the spool read calculations, with spool write IOs/sec being calculated as spool IOs/sec - spool read IOs/sec.

Measurement Discussion: When there are sufficient connections, maximum SMTP throughput is almost entirely determined by the combined effects of four serializing factors: waiting to use the CPU, using the CPU, doing I/O (all synchronous) to the SMTP 191 minidisk, and doing synchronous spool I/O. Each of these factors is mutually exclusive within the SMTP virtual machine. For example, when SMTP is waiting for minidisk I/O to complete, it is not waiting for the CPU, using the CPU, or doing synchronous spool I/O. Waiting for page read I/O can also be a serializing factor but there was no paging during any of the measurement intervals.

We can think of the total capacity of an SMTP server virtual machine as being 1.00 SMTP elapsed seconds per second. Each of the above four components uses up a certain amount of this capacity. This capacity utilization can be calculated from the measured data for all factors except waiting for CPU. These calculations are shown in the results tables. In the case of synchronous spool I/O, the calculations assume that all spool write I/Os are synchronous while, in reality, some spool write I/Os are asynchronous. As such, the synchronous spool I/O calculation represents an upper bound.

The waiting for CPU component can be estimated from the user state data provided in the domain 4 record 4 CP monitor record and shown as the "Waiting on CPU" column in the USER_STATES VMPRF report. This information, collected only for certain runs, is the basis for the "SMTP waiting for CPU-sec/sec" estimates shown in the results tables.

For the send cases, the sum of the four serializing factors (serialized SMTP-sec/sec) is close to 1.00, indicating that these factors are sufficient to account for the observed maximum throughputs. For the receive cases, the sum of the four serializing factors adds up to only 0.79 to 0.83. This indicates that although they do account for most of the observed maximum throughput, there may be one or more additional factors involved as well.

The results show that on the measured configuration, most of the throughput gains arose from the large decrease in I/Os to the SMTP 191 minidisk. The amount of improvement due to this change will be smaller on systems having shorter DASD access times for the SMTP 191 minidisk. This would normally be the case, for example, in cases where write caching is in effect. The send cases also derived some improvement from the use of *SPL to read the spool files asynchronously. The amount of improvement due to this change would be larger on systems with longer spool access times.

If much of the note traffic is between two systems, the connection serialization described earlier can limit throughput.

> Send example: Run SHSR4053 showed a maximum throughput of 3.86 notes sent per second. Four different receiving systems were used for this run. A corresponding measurement with just one receiving system (SHSR1021, not shown), yielded a throughput of 1.73 notes sent per second.
>
> Receive example: Run SHRR3023 showed a maximum throughput of 2.95 notes received per second. Three different sending systems were used for this run. A corresponding measurement with just one sending system (SHRR1041, not shown), yielded a throughput of 2.11 notes received per second.

Equivalent throughput results were obtained when running SMTP MODULE at the 310 level on TCP/IP 2.4 as compared to running all of the TCP/IP code at the 310 level.

The significant NAMESRV CPU usage illustrates that the resolution of host names to internet addresses is an important part of SMTP processing. SMTP′s name resolution requests are asynchronous, so this does not serialize SMTP. However, this does affect throughput per connection and overall efficiency. Because of this, we recommend that you have a domain name server defined as part of the TCP/IP configuration when significant SMTP traffic is anticipated. A caching-only domain name server should be sufficient.

OfficeVision 1.4.0

This section summarizes the results from a pair of measurements obtained to observe the effects of migrating from OfficeVision 1.3.0 to OfficeVision 1.4.0 for the IBM Office Benchmark (IOB). Both measurements were made using VM/ESA 2.2.0 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*	*PAGE*	*SPOOL*	*- Number of Volumes - TDSK*	*User*	*Server*	*System*
3390-3	RAMAC 2	4		8		16		
3390-2	3990-2	4	6		6			
3390-2	3990-3	2						2 R
3390-2	3990-3	4	10		2		7 R	

Note: *R* next to the DASD counts means basic cache enabled; *W* means DASD fast write (and basic cache) enabled. RAMAC 2 refers to the RAMAC 2 Array Subsystem with 256MB cache and drawers in 3390-3 format. With the RAMAC 2 Array Subsystem, read and write caching are always enabled.

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: Overall performance decreased slightly, mostly as a result of changes required for 31 bit support. External response time (AVG LAST(T)) increased by 5%, while internal throughput (ITR(H)) decreased by 1.1%.

The 31 bit support means that the OFSSEG shared segment can now be placed above the 16MB line, freeing space below the line that can be used for other purposes. The EPUYSSEG shared segment (contains the OV/VM mailbox manager code) could already be placed above the line.

This comparison was made using the uncompiled version of the REXX execs provided by OfficeVision. (OfficeVision/VM provides both uncompiled and compiled forms for most of its REXX execs.) A significant percentage of the CPU usage increase is from increased time required to interpret these execs. Therefore, the precentage increase in CPU usage relative to OfficeVision 1.3.0 should be appreciably smaller when using the compiled version of these execs.

Measurements from the *VM/ESA 2.2.0 Performance Report* indicate that OfficeVision 1.2.0 and 1.3.0 have equivalent performance. Therefore, the results shown here also apply to migrations from OfficeVision 1.2.0 to 1.4.0.

Table 25 (Page 1 of 2). Migration from OfficeVision 1.3.0 to OfficeVision 1.4.0 using VM/ESA 2.2.0 on a 9121-480

OV/VM Release **VM/ESA Release** **Run ID**	**1.3.0** **2.2.0** **L29V2102**	**1.4.0** **2.2.0** **L29V2103**	**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	0.048	0.049	0.001	2.08%
NONTRIV INT	0.565	0.587	0.022	3.89%
TOT INT	0.544	0.564	0.020	3.68%
TOT INT ADJ	0.507	0.526	0.018	3.65%
AVG FIRST (T)	0.374	0.392	0.018	4.95%
AVG LAST (T)	0.528	0.555	0.027	5.11%
Throughput				
AVG THINK (T)	46.50	46.40	-0.09	-0.19%
ETR	33.17	33.14	-0.03	-0.09%
ETR (T)	35.58	35.56	-0.02	-0.06%
ETR RATIO	0.932	0.932	0.000	-0.03%
ITR (H)	39.68	39.24	-0.44	-1.10%
ITR	18.52	18.31	-0.21	-1.15%
EMUL ITR	32.20	31.69	-0.51	-1.58%
ITRR (H)	1.000	0.989	-0.011	-1.10%
ITRR	1.000	0.988	-0.012	-1.15%
Proc. Usage				
PBT/CMD (H)	50.402	50.962	0.560	1.11%
PBT/CMD	50.310	50.902	0.593	1.18%
CP/CMD (H)	22.923	23.045	0.122	0.53%
CP/CMD	21.360	21.373	0.013	0.06%
EMUL/CMD (H)	27.479	27.917	0.438	1.59%
EMUL/CMD	28.949	29.529	0.580	2.00%
Processor Util.				
TOTAL (H)	179.33	181.21	1.88	1.05%
TOTAL	179.00	181.00	2.00	1.12%
UTIL/PROC (H)	89.67	90.61	0.94	1.05%
UTIL/PROC	89.50	90.50	1.00	1.12%
TOTAL EMUL (H)	97.77	99.27	1.50	1.53%
TOTAL EMUL	103.00	105.00	2.00	1.94%
MASTER TOTAL (H)	90.21	90.31	0.10	0.11%
MASTER TOTAL	90.00	90.00	0.00	0.00%
MASTER EMUL (H)	38.81	39.70	0.89	2.29%
MASTER EMUL	41.00	42.00	1.00	2.44%
TVR(H)	1.83	1.83	-0.01	-0.48%
TVR	1.74	1.72	-0.01	-0.81%
Storage				
NUCLEUS SIZE (V)	2388KB	2388KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	79	81	2	2.53%
PGBLPGS	55068	55050	-18	-0.03%
PGBLPGS/USER	26.2	26.2	0.0	-0.03%
TOT PAGES/USER (V)	141	144	3	2.13%
FREEPGS	6135	6136	1	0.02%
FREE UTIL	1.00	1.00	0.00	-0.02%
SHRPGS	1534	1552	18	1.17%

Table 25 (Page 2 of 2). Migration from OfficeVision 1.3.0 to OfficeVision 1.4.0 using VM/ESA 2.2.0 on a 9121-480

OV/VM Release **VM/ESA Release** **Run ID**	**1.3.0** **2.2.0** **L29V2102**	**1.4.0** **2.2.0** **L29V2103**	**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	607	627	20	3.29%
WRITES/SEC	551	570	19	3.45%
PAGE/CMD	32.547	33.663	1.117	3.43%
PAGE IO RATE (V)	172.500	178.100	5.600	3.25%
PAGE IO/CMD (V)	4.848	5.009	0.160	3.31%
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	21.276	21.345	0.069	0.32%
Queues				
DISPATCH LIST	35.06	34.93	-0.13	-0.37%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	856	852	-4	-0.47%
VIO/CMD	24.059	23.961	-0.098	-0.41%
RIO RATE (V)	526	530	4	0.76%
RIO/CMD (V)	14.784	14.905	0.121	0.82%
NONPAGE RIO/CMD (V)	9.935	9.896	-0.039	-0.39%
DASD RESP TIME (V)	17.100	17.000	-0.100	-0.58%
MDC REAL SIZE (MB)	53.4	53.4	0.0	-0.05%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	286	285	-1	-0.35%
MDC WRITES (I/Os)	56	55	-1	-1.79%
MDC AVOID	263	263	0	0.00%
MDC HIT RATIO	0.92	0.92	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	21.053	20.922	-0.131	-0.62%
DIAG/CMD	89.898	89.393	-0.505	-0.56%
DIAG 04/CMD	4.400	4.412	0.013	0.28%
DIAG 08/CMD	12.011	11.925	-0.086	-0.72%
DIAG 0C/CMD	4.989	4.981	-0.009	-0.17%
DIAG 14/CMD	1.632	1.624	-0.008	-0.49%
DIAG 58/CMD	2.060	2.060	0.000	0.02%
DIAG 98/CMD	1.435	1.410	-0.025	-1.73%
DIAG A4/CMD	10.730	10.705	-0.025	-0.23%
DIAG A8/CMD	6.334	6.308	-0.026	-0.40%
DIAG 214/CMD	33.727	33.509	-0.218	-0.65%
DIAG 270/CMD	0.667	0.718	0.051	7.67%
SIE/CMD	132.913	133.049	0.137	0.10%
SIE INTCPT/CMD	90.381	90.474	0.093	0.10%
FREE TOTL/CMD	146.909	146.605	-0.305	-0.21%
VTAM Machines				
WKSET (V)	890	890	0	0.00%
TOT CPU/CMD (V)	5.1684	5.1871	0.0187	0.36%
CP CPU/CMD (V)	2.0143	2.0311	0.0168	0.83%
VIRT CPU/CMD (V)	3.1541	3.1560	0.0019	0.06%
DIAG 98/CMD (V)	1.436	1.411	-0.025	-1.75%

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[2]. 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*	*PAGE*	*SPOOL*	*- Number of Volumes - 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* next to the DASD counts means basic cache enabled; *W* means DASD fast write (and basic cache) enabled.

Communications:

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

2 This information first appeared in an update to the *VM/ESA Version 2 Release 2.0 Performance Report*. It is repeated here for those who read the original report but may have missed the update.

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%. Table 26 compares this measurement to an equivalent measurement where each user′s A-filemode was a minidisk. Both measurements were made using OfficeVision 1.3.0.[3]

3 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 26 (Page 1 of 3). OfficeVision performance with an SFS A-directory				
Filemode A **OV/VM Release** **VM/ESA Release** **Run ID**	**Minidisk** **1.3.0** **2.2.0** **L29V2101**	**SFS** **1.3.0** **2.2.0** **L29V1500**	**Difference**	**%Difference**
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 26 (Page 2 of 3). OfficeVision performance with an SFS A-directory

Filemode A OV/VM Release VM/ESA Release Run ID	Minidisk 1.3.0 2.2.0 L29V2101	SFS 1.3.0 2.2.0 L29V1500	Difference	%Difference
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 26 (Page 3 of 3). OfficeVision performance with an SFS A-directory				
Filemode A **OV/VM Release** **VM/ESA Release** **Run ID**	**Minidisk** **1.3.0** **2.2.0** **L29V2101**	**SFS** **1.3.0** **2.2.0** **L29V1500**	**Difference**	**%Difference**
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 6 on page 26 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 2.

<-----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 2. 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.[4]

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.

4 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.

Appendix A. Workloads

The workloads that were used to evaluate VM/ESA 2.3.0 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 27 shows the search-order characteristics of the two environments used for measurements discussed in this document.

Table 27. 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 94 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 V1R2M0
- 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 28 on page 80.

Table 28. 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 29 with their defined use factor.

Table 29. 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.

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:

$$\frac{1}{1000} \times \sum_{f \in \text{filepools}} \frac{\mathit{Agent\_Holding\_Time}_f}{\text{SFSTIME}_f}$$

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:

$$\frac{1}{1000} \times \sum_{f \in \text{filepools}} \frac{\mathit{Filpool\_Request\_Service\_Time}_f}{\text{SFSTIME}_f}$$

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:

$$\frac{\text{Agents In Call}}{\sum_{f \in \text{filepools}} \frac{\mathit{Total\_Filepool\_Requests}_f}{\text{SFSTIME}_f}}$$

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:

$$\frac{1}{\text{ETR (T)}} \times \sum_{t \in \text{TPNS machines}} \frac{\mathit{First\_Response}_t \times \mathit{Total\_Transmits}_t}{\mathit{TPNS\_Time}_t}$$

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:

$$\frac{1}{\text{ETR (T)}} \times \sum_{t \in \text{TPNS machines}} \frac{\mathit{Last\_Response}_t \times \mathit{Total\_Transmits}_t}{\mathit{TPNS\_Time}_t}$$

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:

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

For the PACE workload:

$$\text{PBT/CMD − EMUL/CMD}$$

CP/CMD (H). See CP/CMD. This is the hardware based measure. This is calculated by:

For 9221 processors:
For the FS7F, FS8F, and VSECICS workloads:

$$\frac{CP\_CPU\_PCT \times \text{TOTAL (H)}}{10 \times \text{ETR (T)}}$$

For the PACE workload:

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

CP_CPU_PCT is taken from the Host CPU Busy line in the CPU Busy/MIPs section of the RE0 report.

For all workloads running on 9121 and 9021 processors:

$$\text{PBT/CMD (H) − EMUL/CMD (H)}$$

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

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

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 time a given virtual machine spends using the CPU. This is Total CPU Seconds column for that virtual machine in VMPRF′s USER_RESOURCE_UTIL report divided by run elapsed time.

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 PAGE RATE (V). The number of DASD page reads per second plus DASD page writes per second that occur in a given virtual machine. This is the DASD Read + Write column in VMPRF′s USER_RESOURCE_UTIL report.

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_0C is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 10/CMD. The number of DIAGNOSE code 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 7C/CMD. The number of DIAGNOSE code X7C instructions used per command. DIAGNOSE code X7C, know as the logical device support facility, is the CP function call that enables a virtual machine to communicate with logical 3270 terminals. It is used by the TCP/IP VM Telnet implementation. This is calculated by:

$$\frac{DIAG\_7C}{\text{RTM\_Time} \times \text{ETR (T)}}$$

DIAG_7C 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:

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

Total_BIO_Request_Time is from the QUERY FILEPOOL STATUS command.

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

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

Total_IO_Requests is from the QUERY FILEPOOL STATUS command.

ISFC. Inter-System Facility for Communications

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

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

For the PACE workload:

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

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

For the FS7F, FS8F, and VSECICS workloads:

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

For the PACE workloads:

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

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

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:

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

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:

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

For the PACE workload:

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

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

This is the sum of all the entries in the SSCH+RSCH column for Page devices listed in the VMPRF DASD System Areas by Type report.

PAGE IO/CMD (V). The number of real SSCH and RSCH instructions issued per command on behalf of system paging. This is calculated by:

$$\frac{\text{PAGE IO RATE (V)}}{\text{ETR (T)}}$$

Path length. See Instruction Path Length

PBT/CMD. For the FS7F, FS8F, and VSECICS workloads, this is the number of milliseconds of processor activity per command. For the PACE workload, this is the number of seconds of processor activity per job. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

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

For the PACE workload:

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

PBT/CMD (H). See PBT/CMD. This is the hardware based measure.

For the FS7F, FS8F, and VSECICS workloads:

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

For the PACE workload:

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

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

PD4. One of the CMS-intensive workloads used in the Large Systems Performance Reference (LSPR) to evaluate relative processor performance.

PGBLPGS. The number of system pageable pages available.

This is from the PPAG field for the total RTM time interval (<-) on the RTM SYSTEM screen.

PGBLPGS/USER. The number of system pageable pages available per user. This is calculated by:

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

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

Privileged Operation. Any instruction that must be run in supervisor state.

PRIVOP/CMD. The number of virtual machine privileged instructions simulated per command or job. This does not include DIAGNOSE instructions. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

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

For the PACE workload:

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

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

PRIVOPS (Privileged Operations). See Privileged Operation.

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

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

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

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

PSU. Product Service Upgrade

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

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

RAID. Redundant array of independent DASD.

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

READS/SEC. The number of pages read per second done for system paging.

This is taken from the NSEC column for the PAGREAD field for the total RTM time interval on the RTM SYSTEM screen.

Real Storage. The amount of real storage used for a particular measurement.

Relative Share. A relative share allocates to a virtual machine a portion of the total system resources minus those resources allocated to virtual machines with an ABSOLUTE share. A virtual machine with a RELATIVE share receives access to system resources that is proportional with respect to other virtual machines with RELATIVE shares.

RESERVE. See SET RESERVED

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

RFC. Request for comments. In the context of this report, an RFC is an online document that describes a TCP/IP standard (proposed or adopted).

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:

$$\sum_{f \in \text{filepools}} \mathit{Rollback\_Requests}_f$$

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:

$$\sum_{f \in \text{filepools}} \mathit{Rollbacks\_Due\_to\_Deadlock}_f$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{\mathit{Sac\_Calls}_f}{\text{SFSTIME}_f}}{\sum_{f \in \text{filepools}} \frac{\mathit{Total\_Filepool\_Requests}_f}{\text{SFSTIME}_f}}$$

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

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

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

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:

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

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:

$$\text{TOT INT} \times \text{ETR RATIO}$$

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:

$$\text{UTIL/PROC (H)} \times \text{PROCESSORS}$$

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

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

TOTAL EMUL. The total emulation state time for all users across all online processors. This indicates the percentage of time the processors are in emulation state.

This comes from the %EM column for all processors for the total RTM interval time (<-..) on the RTM CPU screen.

TOTAL EMUL (H). The total emulation state time for all users across all online processors. This indicates the percentage of time the processors are in emulation state. This is calculated by:

For 9221 processors, this comes from the SIE CPU Busy / Total CPU Busy (PCT) line in the RE0 report.

For 9121 and 9021 processors, this comes from the %CPU column for the GUES-ALL line of the REPORT file times the number of processors.

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

TPNS. Teleprocessing Network Simulator. A licensed program terminal and network simulation tool that provides system performance and response time information.

Transaction. A user/system interaction as counted by CP. For a single-user virtual machine a transaction should roughly correspond to a command. It does not include network or transmission delays and may include false transactions. False transactions can be those that wait for an external event, causing them to be counted as multiple transactions, or those that process more than one command without dropping from queue, causing multiple transactions to be counted as one.

TRACE TABLE (V). The size in kilobytes of the CP trace table.

This is the value of the <K bytes> column on the Trace Table line in the VMPRF System Configuration Report.

Transaction (T). This is the interval from the time the command is issued until the last receive prior to the next send. This includes clear screens as a result of an intervening MORE... or HOLDING condition.

TRIV INT. Trivial Internal Response Time in seconds. The average response time for transactions that complete with one and only one drop from Q1 and no drops from Q0, Q2, and Q3.

This is from TOTALTTM for the TRIV field on the RTM TRANSACT screen.

TVR. Total to Virtual Ratio. This is the ratio of total processor utilization to virtual processor utilization. This is calculated by:

$$\frac{\text{TOTAL}}{\text{TOTAL EMUL}}$$

TVR (H). See TVR. Total to Virtual Ratio measured by the hardware monitor. This is calculated by:

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

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:

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

For the PACE workload:

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