Virtual Machine/
Enterprise Systems Architecture
Performance Report
Version 2 Release 4

July 6, 1999

VM Performance Group

IBM Corporation
Endicott, New York
ernsberw@us.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 2.4.0 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 2.4.0. See the IBM Programming Announcement for VM/ESA 2.4.0 for more information about what publications are considered to be product documentation.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

AIX
ACF/VTAM
CICS
DisplayWrite
ES/9000
FICON
IBM
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

Tivoli ADSM is a trademark of Tivoli Systems Inc. in the United States, or other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names may be trademarks or service marks of others.

Acknowledgements

The following people contributed to this report:

Bill Bitner
Wes Ernsberger
Bill Guzior

Abstract

The *VM/ESA Version 2 Release 4.0 Performance Report* summarizes the performance evaluation of VM/ESA 2.4.0, including TCP/IP Function Level 320. Measurements were obtained for CMS-intensive, VSE guest, Telnet, FTP, and NFS environments on various ES/9000 processors. Discussion covers the performance changes in VM/ESA 2.4.0, the performance effects of migrating from VM/ESA 2.3.0 to VM/ESA 2.4.0, the performance effects of migrating from TCP/IP 310 to TCP/IP 320, and additional evaluations.

Referenced Publications

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

- *VM/ESA: CMS File Pool Planning, Administration, and Operation,* SC24-5751
- *TCP/IP Function Level 320 Planning and Customization,* SC24-5847

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.ibm.com/s390/vm/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,* see http://www.ibm.com/s390/vm/perf/docs/
- *VM/ESA Version 2 Release 3.0 Performance Report,* see http://www.ibm.com/s390/vm/perf/docs/

Much additional VM/ESA performance information is available on the VM/ESA performance page at http://www.ibm.com/s390/vm/perf/.

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

Summary of Key Findings

This report summarizes the performance evaluation of VM/ESA Version 2 Release 4.0, including TCP/IP Function Level 320. Measurements were obtained for the CMS-intensive, VSE guest, Telnet, FTP, and NFS 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 2.4.0 includes a number of performance enhancements (page 3) and changes that affect VM/ESA performance management (page 5):

- Performance Improvements
  - VM/ESA
    - CMS Multitasking Improvements
    - Improved CMS EXECs, Xedit Macros, and PIPELINE Stages
    - Improved SFS Locking Characteristics
  - TCP/IP
    - Improved NFS Performance
    - TCP Layer Pathlength Reductions
    - Improved Window Scaling
    - Maximum Large Envelope Size Increased
    - FTP Client Pathlength Reduction
- Performance Management Changes
  - Improved LIMITHARD Share Capping Option
  - Monitor Enhancements
  - Effects on Accounting Data
  - VM Performance Products

Migration from VM/ESA 2.3.0 and TCP/IP 310

Benchmark measurements show the following performance results for VM/ESA 2.4.0 relative to VM/ESA 2.3.0:

CMS-intensive Overall performance was similar to VM/ESA 2.3.0. The internal throughput rate (ITR) decreased by an average 0.6% while response times were equivalent (page 11).

VSE guest Performance was equivalent for both the V=R and V=V environments (page 22).

Benchmark measurements show the following performance results for TCP/IP FL 320 relative to TCP/IP FL 310:

Telnet Performance was equivalent (page 26).

FTP Performance of the “get” cases improved while performance of the “put” cases was unchanged. On average, FTP get throughput improved by 2% while total CPU usage improved by 6% (page 29).

NFS Elapsed time improvements of up to 83% and processor usage improvements of up to 87% have been measured for NFS workloads. NFS access to BFS files improved the most. The performance of NFS access to minidisk, SFS, and BFS files is now much more similar than it was with TCP/IP FL 310 (page 30).

Additional Evaluations

Measurement results indicate that the actual processor utilization observed for a virtual machine now tracks more closely to the processor utilization cap that was requested for it by use of the LIMITHARD option of the SET SHARE command (page 47).

Measurement results demonstrate that the Reusable Server Kernel can sustain high levels of throughput while retaining good response times (page 48).

Tivoli ADSM Version 3 offers substantially improved performance relative to Version 2. For backup, throughput improvements ranging from 7% to 188% were observed, while VM processor and I/O usage improved by 17% to 66% and by 77% to 87% respectively. Restore improvements were smaller. Throughput improvements ranged from -3% to 15%, while VM processor and I/O usage improved by 21% to 52% and by 0% to 47% respectively (page 51).

Changes That Affect Performance

This chapter contains descriptions of various changes to VM/ESA 2.4.0 that affect performance. It is divided into two sections — "Performance Improvements" and "Performance Management." This information is also available at http://www.ibm.com/s390/vm/perf, along with corresponding information for previous releases.

Throughout the rest of this report, 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 Improvements

The following items improve the performance of VM/ESA.

- VM/ESA
  - CMS Multitasking Improvements
  - Improved CMS EXECs, Xedit Macros, and PIPELINE Stages
  - Improved SFS Locking Characteristics
- TCP/IP
  - Improved NFS Performance
  - TCP Layer Pathlength Reductions
  - Improved Window Scaling
  - Maximum Large Envelope Size Increased
  - FTP Client Pathlength Reduction

CMS Multitasking Improvements

The pathlengths associated with the thread creation, event signalling, and Posix signal handling CMS multitasking functions have been significantly reduced in VM/ESA 2.4.0. This improvement is especially applicable to situations that involve access to byte file system (BFS) files. Examples: Java initialization processor usage was reduced by 10%. NFS server processor usage for reading a large BFS file decreased by 27% (see page 33).

Improved CMS EXECs, Xedit Macros, and PIPELINE Stages

The CMS Pipelines DATECONVERT stage is implemented in REXX. VM/ESA 2.3.0 APAR VM61673 provided this stage in compiled form for improved performance. This improvement has now been integrated into VM/ESA 2.4.0.

PEEK is often used to view the beginning of a large reader file (PEEK defaults to reading the first 200 records). With VM/ESA 2.3.0, PEEK would read the entire contents of a reader file into memory if it was in NETDATA (used by SENDFILE) or disk dump format. With VM/ESA 2.4.0, PEEK has been changed so that only the requested records are kept in memory. For large reader files in NETDATA or disk dump format, this can result in greatly improved responsiveness and reduced requirements for virtual storage and processor time.

The SORT Xedit macro, formerly written in EXEC2, has been rewritten in REXX, compiled, and added to the CMSINST shared segment. This reduces processor, I/O, and real storage usage. The measured processor usage improvements were usually in the 2% to 4% range.

The SDIR Xedit macro has been rewritten for improved serviceability and performance, resulting in a measured 38% processor usage improvement. SDIR is used by the FILELIST function to sort by SFS subdirectory.

Processor usage of SENDFILE when using the UFTSYNC option has been significantly reduced.

The RECEIVE EXEC was added to the CMSINST saved segment, eliminating loading time and reducing real storage requirements for CMS environments that use this function.

Improved SFS Locking Characteristics

Three SFS performance APARs have been incorporated into VM/ESA 2.4.0. They all deal with scenarios where other SFS users appear to be locked out while a particular task for another user is in progress. VM61547 and VM62008 deals with the task of deleting large files. The impact is proportional to file size and is not really noticeable for files under 512 Kb. VM62086 addresses a different scenario where a large number of file changes are made without a commit being issued.

Improved NFS Performance

Elapsed time improvements of up to 83% and processor usage improvements of up to 87% have been measured for NFS workloads. These improvements result from the combined effects of performance enhancements that were added to NFS through service and in TCP/IP Function Level 320, support for the NFS version 3 protocol, and the CMS multitasking improvements that are included in VM/ESA 2.4.0. NFS access to BFS and SFS files improved the most but NFS access to minidisk files also improved in some cases. See "NFS" on page 30 for details.

TCP Layer Pathlength Reductions

Two improvements were made to the TCP layer of the protocol stack that serve to reduce the mainline pathlength for handling incoming messages. The first one, sometimes referred to as "TCP header prediction," optimizes the implementation for the normal case where the incoming segments are all present and are received in the order sent. The other improvement is the creation of a lookaside buffer for fast access to the TCP control block. Primarily due to these improvements, processor usage in the TCPIP stack virtual machine decreased by 0.5% for the Telnet measurement (see "Telnet" on page 26) and by an average of 4% for the 6 FTP measurements (see "FTP" on page 29). Large data transfers tend to experience the largest percentage improvements.

Improved Window Scaling

APAR PQ18391 to TCP/IP Function Level 310 is now included in TCP/IP FL 320. Previously, there was a hardcoded limit of 20 segments that could be outstanding (unacknowledged) at any given time. This was a constraint whenever the window size exceeded 20 times the maximum segment size. This 20 segment limit has now been removed, allowing the full benefits of TCP window scaling (RFC 1323) to be realized.

Maximum Large Envelope Size Increased

The maximum large envelope size supported by TCP/IP VM has been increased from 32768 bytes (32 Kb) to 65535 (64 Kb - 1). The TCP/IP stack machine uses large envelopes to hold UDP datagrams and to hold IP datagram fragments during reassembly when they do not fit into small (2048 bytes) envelopes. Large envelope size is specified on the LARGEENVELOPEPOOLSIZE statement in the *node_name* TCPIP (or PROFILE TCPIP) configuration file.

A large envelope size greater than 32 Kb may improve performance in certain cases. One example is when a key application can send datagrams that exceed 32 Kb. Another example is when it is advantageous to be able to specify an MTU size that exceeds 32 Kb as might be the case, for example, for CTC devices. TCP/IP VM requires that the MTU size (specified on the GATEWAY statement) not exceed the large envelope size.

FTP Client Pathlength Reduction

A change has been made to FTP MODULE that reduces pathlength in the VM client virtual machine for the case of binary file transfer from a foreign host to the VM system. This changed was observed to decrease processor usage in the client virtual machine by 9% for binary get of a 2 Mb file and by 31% for binary get of a 24 Kb file (see "FTP" on page 29 for details). This change is implemented by TCP/IP FL 320 APAR PQ28148.

Performance Management

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

- Improved LIMITHARD Share Capping Option
- Monitor Enhancements
- Effects on Accounting Data
- VM Performance Products

Improved LIMITHARD Share Capping Option

The LIMITHARD option on the SET SHARE command can be used to limit the percentage of total system processing capacity that can be consumed by a given virtual machine. In prior releases, the observed percentage would tend to be lower than the requested percentage, especially at low total system processor utilizations. With VM/ESA 2.4.0, the LIMITHARD implementation has been changed so that the actual percentage tracks more closely to the requested percentage. See "LIMITHARD Share Capping Option" on page 47 for measurement results and additional discussion.

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 2.4.0. The major changes in the CP monitor data were extended channel measurement support, inclusion of data for new DASD, and some miscellaneous changes.

With the addition of support for FICON channels, some processors have also added the Extended Channel Path Measurement Facility. The original Channel Path Measurement Facility is still supported and is used if the extended facility is not available. A separate record, domain 0 record 20, is created for each channel when the extended facilities are available. For the new FICON channels, this new record will include channel utilization and number of bytes read and written over the channel for the physical channel and also from a logical view when VM/ESA is running in a logical partition. Comments were changed in other channel utilization records to reference this new record where appropriate.

Changes were also made in support of features provided by the IBM Enterprise Storage Server (ESS). The ESS DASD subsystem allows for parallel access volumes where a single physical volume can appear as multiple volumes to a S/390 system. This allows multiple I/Os to be started to the same device in parallel. The Device Configuration (domain 1 record 6) and Vary On Device (domain 6 record 1) were changed to indicate whether the specified volume is a parallel access volume, and if so whether it is the base volume or an alias volume. Parallel access volumes are typically configured prior to system IPLs. However, the State Change event record (domain 6 record 20) was added to record the dynamic creation of an alias volume.

The System Configuration Data (domain 1 record 4) was enhanced to indicate the source for Subchannel Device Active Only time (as reported in domain 6 record 3). On some processors this data is generated by the hardware, while on others software generates it. Also, an indication of whether the Channel Path Measurement Facility or Extended Channel Path Measurement Facility is available on the processor has been added to the System Configuration Data record.

The APPLDATA domain data contributed by the TCP/IP stack virtual machine has also been enhanced. The changes include the following.

- In order to record larger data transfer amounts, new fields were added to the TCB Close (type 2) and UCB Close (type 7) event records. This change was introduced via APAR PQ16942 to FL310 and rolled into the base for FL320.
- Support was added in FL320 to handle certain denial of service attacks. Discarded packets associated with these attacks are counted in the MIB record (type 0).
- The IP address for a connection was added to the TCB Open (type 1) and Close (type 2) records.
- Window scaling information was added to the TCB Close (type 2) record.
- The FL320 stack was optimized to predict TCP headers. A counter was added to the TCB Close (type 2) record to indicate the number of times headers were predicted correctly.
- The TCP/IP Pool Limit Record (type 3) now includes additional information for the new dynamic pools and for virtual storage usage in general for the stack machine. Similarly, the TCP/IP Pool Size Record (type 4) now includes the new segment acknowledge pool.

Effects on Accounting Data

None of the VM/ESA 2.4.0 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 2.4.0 provided by VMPRF, RTM/ESA, FCON/ESA, and VMPAF.

VM Performance Reporting Facility 1.2.1 (VMPRF) will run on VM/ESA 2.4.0 with the same support as VM/ESA 2.3.0. The latest service is recommended.

Realtime Monitor VM/ESA 1.5.2 (RTM/ESA) will run on VM/ESA 2.4.0 with the same support as VM/ESA 2.3.0. RTM must be recompiled using the VM/ESA 2.4.0 libraries. Follow the ″REBUILD″ section of the RTM Program Directory.

FCON/ESA Version 3.1.xx will run on VM/ESA 2.4.0 with the same support as for VM/ESA 2.3.0. The next release, FCON/ESA V.3.2.00, planned to become available in the third quarter of 1999, will also include support for FICON channels and some of the additional TCP/IP FL 320 data, in addition to other enhancements that are not VM/ESA release dependent.

Performance Analysis Facility/VM 1.1.3 (VMPAF) will run on VM/ESA 2.4.0 with the same support as VM/ESA 2.3.0.

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.
- 9672-R55

Software

A pre-GA (General Availability) level of VM/ESA 2.4.0 was used for the measurements in this report.

Other VM/ESA releases were measured for this report. VM/ESA 2.3.0 was at the GA+first-RSU (Recommended Service Upgrade) level. The service that was part of VM/ESA 2.3.0 after the first RSU level and integrated into VM/ESA 2.4.0 can account for some of the performance differences between VM/ESA 2.3.0 and VM/ESA 2.4.0.

See the appropriate workload section in Appendix A, "Workloads" on page 58 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 58.
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.

   - CMS block size: The block size of the CMS minidisks.

   - Virtual Machines: The virtual machines used in the measurement.

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

4. Measurement Discussion: This contains an analysis of the performance data in the table and gives the overall performance findings.

5. Measurement Data: This contains the table of performance results. These data were obtained or derived from the tools listed in “Tools Description” on page 10.

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

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

Tools Description

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

Licensed Programs:

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

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

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

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

Internal Tools:

FSTTAPE Reduces hardware monitor data for the 9121 processors.

Hardware Monitor Collects processor event and timing data.

REDFP Consolidates the QUERY FILEPOOL STATUS data from SFS measurements.

Migration from VM/ESA 2.3.0 and TCP/IP FL 310

This chapter examines the performance effects of migrating from VM/ESA 2.3.0 to VM/ESA 2.4.0 and from TCP/IP Function Level 310 to TCP/IP Function Level 320. The following environments were measured: CMS-intensive, VSE guest, Telnet, FTP, and NFS.

CMS-Intensive

The VM/ESA 2.4.0 performance results are similar to the corresponding VM/ESA 2.3.0 results for the three measured CMS-intensive environments. ITR decreases ranging from 0.6% to 0.7% and external response time changes ranging from a 2% decrease to a 6% increase were observed.

Although VM/ESA 2.4.0 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.

All measurements were done using the FS8F0R workload. See "CMS-Intensive (FS8F)" on page 58 for a description.

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		6		32		
3390-2	3990-3	4	6		2			2 R
3390-2	3990-2	4	6		2			
3390-2	3990-2	4	10					

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 2.4.0 shows slightly reduced performance relative to VM/ESA 2.3.0 for this workload and system configuration. External response time (AVG LAST(T)) increased by 6%, while the internal throughput rate (ITR(H)) decreased by 0.7%. DASD paging increased slightly due to a 2-page increase in the average working set size.

Table 1 (Page 1 of 3). Minidisk-only CMS-intensive migration from VM/ESA 2.3.0 on the 9121-742

Release Run ID	**2.3.0 S4AE5805**	**2.4.0 S4BE5801**	**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	0.112	0.112	0.000	0.00%
NONTRIV INT	0.359	0.367	0.008	2.23%
TOT INT	0.258	0.260	0.002	0.78%
TOT INT ADJ	0.257	0.262	0.006	2.20%
AVG FIRST (T)	0.308	0.329	0.021	6.68%
AVG LAST (T)	0.395	0.419	0.024	6.05%
Throughput				
AVG THINK (T)	26.20	26.22	0.02	0.08%
ETR	202.44	205.16	2.72	1.34%
ETR (T)	203.46	203.31	-0.15	-0.07%
ETR RATIO	0.995	1.009	0.014	1.42%
ITR (H)	222.71	221.20	-1.51	-0.68%
ITR	55.47	55.85	0.39	0.70%
EMUL ITR	87.90	88.99	1.09	1.24%
ITRR (H)	1.000	0.993	-0.007	-0.68%
ITRR	1.000	1.007	0.007	0.70%
Proc. Usage				
PBT/CMD (H)	17.961	18.083	0.122	0.68%
PBT/CMD	17.940	18.101	0.161	0.90%
CP/CMD (H)	7.048	7.167	0.119	1.69%
CP/CMD	6.635	6.739	0.103	1.56%
EMUL/CMD (H)	10.912	10.916	0.003	0.03%
EMUL/CMD	11.305	11.362	0.057	0.51%

Table 1 (Page 2 of 3). Minidisk-only CMS-intensive migration from VM/ESA 2.3.0 on the 9121-742

Release **Run ID**	**2.3.0** **S4AE5805**	**2.4.0** **S4BE5801**	**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)	365.42	367.64	2.22	0.61%
TOTAL	365.00	368.00	3.00	0.82%
UTIL/PROC (H)	91.36	91.91	0.56	0.61%
UTIL/PROC	91.25	92.00	0.75	0.82%
TOTAL EMUL (H)	222.02	221.92	-0.10	-0.04%
TOTAL EMUL	230.00	231.00	1.00	0.43%
MASTER TOTAL (H)	93.13	93.68	0.56	0.60%
MASTER TOTAL	93.00	94.00	1.00	1.08%
MASTER EMUL (H)	35.06	34.67	-0.38	-1.10%
MASTER EMUL	37.00	36.00	-1.00	-2.70%
TVR(H)	1.65	1.66	0.01	0.65%
TVR	1.59	1.59	0.01	0.39%
Storage				
NUCLEUS SIZE (V)	2452KB	2476KB	24KB	0.98%
TRACE TABLE (V)	650KB	650KB	0KB	0.00%
WKSET (V)	78	80	2	2.56%
PGBLPGS	230K	230K	0K	0.00%
PGBLPGS/USER	39.7	39.7	0.0	0.00%
TOT PAGES/USER (V)	193	193	0	0.00%
FREEPGS	17772	17765	-7	-0.04%
FREE UTIL	0.91	0.91	0.00	0.04%
SHRPGS	2165	2183	18	0.83%
Paging				
READS/SEC	1236	1310	74	5.99%
WRITES/SEC	871	930	59	6.77%
PAGE/CMD	10.356	11.018	0.662	6.39%
PAGE IO RATE (V)	350.300	369.700	19.400	5.54%
PAGE IO/CMD (V)	1.722	1.818	0.097	5.61%
XSTOR IN/SEC	549	516	-33	-6.01%
XSTOR OUT/SEC	1612	1641	29	1.80%
XSTOR/CMD	10.622	10.610	-0.012	-0.11%
FAST CLR/CMD	8.474	8.657	0.183	2.16%
Queues				
DISPATCH LIST	113.70	120.44	6.74	5.93%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	1839	1836	-3	-0.16%
VIO/CMD	9.039	9.031	-0.008	-0.09%
RIO RATE (V)	744	764	20	2.69%
RIO/CMD (V)	3.657	3.758	0.101	2.76%
NONPAGE RIO/CMD (V)	1.935	1.939	0.004	0.23%
DASD RESP TIME (V)	21.400	21.600	0.200	0.93%
MDC REAL SIZE (MB)	32.1	30.3	-1.8	-5.69%
MDC XSTOR SIZE (MB)	63.5	63.6	0.2	0.25%
MDC READS (I/Os)	568	566	-2	-0.35%
MDC WRITES (I/Os)	28	28	0	0.00%
MDC AVOID	529	524	-5	-0.95%
MDC HIT RATIO	0.93	0.92	-0.01	-1.08%

Table 1 (Page 3 of 3). Minidisk-only CMS-intensive migration from VM/ESA 2.3.0 on the 9121-742

Release **Run ID**	**2.3.0** **S4AE5805**	**2.4.0** **S4BE5801**	**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	20.214	20.191	-0.023	-0.11%
DIAG/CMD	23.923	24.866	0.942	3.94%
DIAG 04/CMD	1.088	1.013	-0.075	-6.93%
DIAG 08/CMD	0.735	0.734	-0.001	-0.09%
DIAG 0C/CMD	0.192	0.192	0.000	-0.14%
DIAG 14/CMD	0.024	0.024	0.000	0.13%
DIAG 58/CMD	1.250	1.249	-0.001	-0.07%
DIAG 98/CMD	0.304	0.298	-0.007	-2.15%
DIAG A4/CMD	3.495	3.468	-0.027	-0.76%
DIAG A8/CMD	2.667	2.673	0.006	0.22%
DIAG 214/CMD	11.851	12.754	0.903	7.62%
DIAG 270/CMD	0.941	0.940	-0.001	-0.09%
SIE/CMD	54.066	54.105	0.039	0.07%
SIE INTCPT/CMD	36.224	36.251	0.026	0.07%
FREE TOTL/CMD	45.125	45.262	0.136	0.30%
VTAM Machines				
WKSET (V)	4129	4140	11	0.27%
TOT CPU/CMD (V)	2.8699	2.8938	0.0239	0.83%
CP CPU/CMD (V)	1.2725	1.2925	0.0200	1.57%
VIRT CPU/CMD (V)	1.5974	1.6013	0.0039	0.24%
DIAG 98/CMD (V)	0.304	0.298	-0.006	-2.08%

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		4		16		
3390-2	3990-2	4	5					
3390-2	3990-3	2	6		1			2 R
3390-2	3990-3	4	5		1			

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)) was equivalent, while internal throughput (ITR(H)) decreased by 0.6%. DASD paging increased slightly due to a 2-page increase in the average working set size. These results are similar to what was observed for the 9121-742 environment (see the previous section).

Migration: CMS-Intensive

Table 2 (Page 1 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.3.0 on the 9121-480

Release **Run ID**	**2.3.0** **L2AE2047**	**2.4.0** **L2BE2042**	**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	0.129	0.133	0.004	3.10%
NONTRIV INT	0.428	0.432	0.004	0.93%
TOT INT	0.326	0.330	0.004	1.23%
TOT INT ADJ	0.294	0.297	0.003	1.03%
AVG FIRST (T)	0.283	0.285	0.002	0.71%
AVG LAST (T)	0.391	0.393	0.001	0.38%
Throughput				
AVG THINK (T)	26.14	26.17	0.03	0.11%
ETR	64.41	64.27	-0.14	-0.22%
ETR (T)	71.40	71.38	-0.02	-0.02%
ETR RATIO	0.902	0.900	-0.002	-0.19%
ITR (H)	79.55	79.07	-0.48	-0.61%
ITR	35.91	35.61	-0.30	-0.83%
EMUL ITR	53.53	53.21	-0.32	-0.59%
ITRR (H)	1.000	0.994	-0.006	-0.61%
ITRR	1.000	0.992	-0.008	-0.83%
Proc. Usage				
PBT/CMD (H)	25.140	25.293	0.153	0.61%
PBT/CMD	25.211	25.357	0.146	0.58%
CP/CMD (H)	8.851	8.965	0.114	1.29%
CP/CMD	8.404	8.406	0.002	0.02%
EMUL/CMD (H)	16.289	16.328	0.039	0.24%
EMUL/CMD	16.807	16.951	0.144	0.86%
Processor Util.				
TOTAL (H)	179.49	180.54	1.05	0.59%
TOTAL	180.00	181.00	1.00	0.56%
UTIL/PROC (H)	89.75	90.27	0.53	0.59%
UTIL/PROC	90.00	90.50	0.50	0.56%
TOTAL EMUL (H)	116.30	116.55	0.25	0.21%
TOTAL EMUL	120.00	121.00	1.00	0.83%
MASTER TOTAL (H)	89.36	89.98	0.62	0.69%
MASTER TOTAL	89.00	90.00	1.00	1.12%
MASTER EMUL (H)	51.12	51.28	0.16	0.32%
MASTER EMUL	53.00	53.00	0.00	0.00%
TVR(H)	1.54	1.55	0.01	0.37%
TVR	1.50	1.50	0.00	-0.28%
Storage				
NUCLEUS SIZE (V)	2452KB	2476KB	24KB	0.98%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	83	85	2	2.41%
PGBLPGS	54139	54135	-4	-0.01%
PGBLPGS/USER	26.5	26.5	0.0	-0.01%
TOT PAGES/USER (V)	172	173	1	0.58%
FREEPGS	6245	6249	4	0.06%
FREE UTIL	0.94	0.94	0.00	-0.06%
SHRPGS	1566	1670	104	6.64%

Table 2 (Page 2 of 2). Minidisk-only CMS-intensive migration from VM/ESA 2.3.0 on the 9121-480

Release **Run ID**	**2.3.0** **L2AE2047**	**2.4.0** **L2BE2042**	**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	728	740	12	1.65%
WRITES/SEC	474	484	10	2.11%
PAGE/CMD	16.835	17.148	0.312	1.85%
PAGE IO RATE (V)	202.500	205.100	2.600	1.28%
PAGE IO/CMD (V)	2.836	2.873	0.037	1.31%
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	8.404	8.588	0.184	2.19%
Queues				
DISPATCH LIST	43.20	43.80	0.60	1.38%
ELIGIBLE LIST	0.02	0.00	-0.02	-100.00%
I/O				
VIO RATE	693	691	-2	-0.29%
VIO/CMD	9.706	9.680	-0.026	-0.27%
RIO RATE (V)	395	398	3	0.76%
RIO/CMD (V)	5.532	5.576	0.043	0.78%
NONPAGE RIO/CMD (V)	2.696	2.702	0.006	0.23%
DASD RESP TIME (V)	19.900	19.800	-0.100	-0.50%
MDC REAL SIZE (MB)	40.6	39.8	-0.7	-1.75%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	199	198	-1	-0.50%
MDC WRITES (I/Os)	9.81	9.84	0.03	0.31%
MDC AVOID	187	186	-1	-0.53%
MDC HIT RATIO	0.93	0.93	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	13.890	13.856	-0.035	-0.25%
DIAG/CMD	25.845	26.820	0.975	3.77%
DIAG 04/CMD	2.321	2.287	-0.034	-1.48%
DIAG 08/CMD	0.737	0.733	-0.004	-0.53%
DIAG 0C/CMD	0.192	0.192	0.000	0.08%
DIAG 14/CMD	0.025	0.024	0.000	-1.25%
DIAG 58/CMD	1.249	1.249	0.000	0.00%
DIAG 98/CMD	0.970	0.963	-0.007	-0.69%
DIAG A4/CMD	3.503	3.488	-0.015	-0.43%
DIAG A8/CMD	2.661	2.658	-0.003	-0.12%
DIAG 214/CMD	11.866	12.764	0.898	7.57%
DIAG 270/CMD	0.940	0.940	0.000	-0.03%
SIE/CMD	52.803	53.152	0.348	0.66%
SIE INTCPT/CMD	34.322	34.017	-0.305	-0.89%
FREE TOTL/CMD	49.064	49.565	0.502	1.02%
VTAM Machines				
WKSET (V)	538	553	15	2.79%
TOT CPU/CMD (V)	3.8673	3.8759	0.0086	0.22%
CP CPU/CMD (V)	1.4473	1.4554	0.0081	0.56%
VIRT CPU/CMD (V)	2.4200	2.4205	0.0005	0.02%
DIAG 98/CMD (V)	0.971	0.963	-0.007	-0.75%

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 2%, while internal throughput (ITR(H)) decreased by 0.6%. This environment did not experience a working set increase. The ITR decrease is the same as observed for the 9121-480 minidisk environment (see the previous section).

Table 3 (Page 1 of 3). SFS CMS-intensive migration from VM/ESA 2.3.0 on the 9121-480

Release **Run ID**	**2.3.0** **L2AS1700**	**2.4.0** **L2BS1701**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1700**	**1700**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Response Time				
TRIV INT	0.126	0.124	-0.002	-1.59%
NONTRIV INT	0.464	0.443	-0.021	-4.53%
TOT INT	0.353	0.337	-0.016	-4.53%
TOT INT ADJ	0.310	0.297	-0.013	-4.35%
AVG FIRST (T)	0.251	0.251	0.000	0.20%
AVG LAST (T)	0.374	0.365	-0.009	-2.41%
Throughput				
AVG THINK (T)	26.15	26.18	0.03	0.11%
ETR	52.41	52.53	0.12	0.23%
ETR (T)	59.65	59.67	0.02	0.04%
ETR RATIO	0.879	0.880	0.002	0.19%
ITR (H)	66.85	66.43	-0.42	-0.63%
ITR	29.38	29.27	-0.11	-0.38%
EMUL ITR	44.60	44.52	-0.08	-0.18%
ITRR (H)	1.000	0.994	-0.006	-0.63%
ITRR	1.000	0.996	-0.004	-0.38%
Proc. Usage				
PBT/CMD (H)	29.917	30.107	0.190	0.63%
PBT/CMD	29.841	30.164	0.323	1.08%
CP/CMD (H)	10.861	10.992	0.131	1.21%
CP/CMD	10.059	10.390	0.331	3.29%
EMUL/CMD (H)	19.056	19.115	0.059	0.31%
EMUL/CMD	19.782	19.774	-0.008	-0.04%
Processor Util.				
TOTAL (H)	178.46	179.66	1.20	0.67%
TOTAL	178.00	180.00	2.00	1.12%
UTIL/PROC (H)	89.23	89.83	0.60	0.67%
UTIL/PROC	89.00	90.00	1.00	1.12%
TOTAL EMUL (H)	113.67	114.06	0.39	0.35%
TOTAL EMUL	118.00	118.00	0.00	0.00%
MASTER TOTAL (H)	89.06	89.60	0.54	0.61%
MASTER TOTAL	89.00	90.00	1.00	1.12%
MASTER EMUL (H)	51.17	51.17	0.00	0.00%
MASTER EMUL	53.00	53.00	0.00	0.00%
TVR(H)	1.57	1.58	0.01	0.33%
TVR	1.51	1.53	0.02	1.12%
Storage				
NUCLEUS SIZE (V)	2452KB	2476KB	24KB	0.98%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	81	81	0	0.00%
PGBLPGS	55413	55403	-10	-0.02%
PGBLPGS/USER	32.6	32.6	0.0	-0.02%
TOT PAGES/USER (V)	159	162	3	1.89%
FREEPGS	5358	5384	26	0.49%
FREE UTIL	0.91	0.90	0.00	-0.48%
SHRPGS	1851	1822	-29	-1.57%

Table 3 (Page 2 of 3). SFS CMS-intensive migration from VM/ESA 2.3.0 on the 9121-480

Release **Run ID**	**2.3.0** **L2AS1700**	**2.4.0** **L2BS1701**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1700**	**1700**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	595	616	21	3.53%
WRITES/SEC	390	394	4	1.03%
PAGE/CMD	16.513	16.925	0.412	2.50%
PAGE IO RATE (V)	154.100	154.600	0.500	0.32%
PAGE IO/CMD (V)	2.583	2.591	0.007	0.29%
XSTOR IN/SEC	0	0	0	na
XSTOR OUT/SEC	0	0	0	na
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	8.164	8.278	0.114	1.40%
Queues				
DISPATCH LIST	43.02	38.77	-4.25	-9.88%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	592	590	-2	-0.34%
VIO/CMD	9.925	9.887	-0.037	-0.38%
RIO RATE (V)	357	356	-1	-0.28%
RIO/CMD (V)	5.985	5.966	-0.019	-0.32%
NONPAGE RIO/CMD (V)	3.402	3.375	-0.026	-0.78%
DASD RESP TIME (V)	17.900	18.000	0.100	0.56%
MDC REAL SIZE (MB)	67.1	66.6	-0.5	-0.79%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	155	155	0	0.00%
MDC WRITES (I/Os)	15	15	0	0.00%
MDC AVOID	129	128	-1	-0.78%
MDC HIT RATIO	0.82	0.82	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	20.688	20.623	-0.065	-0.32%
DIAG/CMD	24.009	25.012	1.003	4.18%
DIAG 04/CMD	2.484	2.532	0.048	1.95%
DIAG 08/CMD	0.739	0.740	0.001	0.16%
DIAG 0C/CMD	0.218	0.214	-0.004	-1.80%
DIAG 14/CMD	0.025	0.025	0.000	0.62%
DIAG 58/CMD	1.250	1.249	-0.001	-0.08%
DIAG 98/CMD	1.222	1.187	-0.035	-2.85%
DIAG A4/CMD	1.905	1.896	-0.010	-0.51%
DIAG A8/CMD	2.486	2.495	0.009	0.37%
DIAG 214/CMD	11.363	12.216	0.854	7.51%
DIAG 270/CMD	0.941	0.942	0.000	0.03%
SIE/CMD	60.788	60.965	0.177	0.29%
SIE INTCPT/CMD	41.336	41.456	0.121	0.29%
FREE TOTL/CMD	52.707	53.056	0.348	0.66%
VTAM Machines				
WKSET (V)	519	492	-27	-5.20%
TOT CPU/CMD (V)	4.0980	4.0871	-0.0109	-0.27%
CP CPU/CMD (V)	1.5274	1.5268	-0.0006	-0.04%
VIRT CPU/CMD (V)	2.5706	2.5602	-0.0104	-0.40%
DIAG 98/CMD (V)	1.222	1.186	-0.035	-2.90%

Table 3 (Page 3 of 3). SFS CMS-intensive migration from VM/ESA 2.3.0 on the 9121-480

Release **Run ID**	**2.3.0** **L2AS1700**	**2.4.0** **L2BS1701**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**00MB**	**00MB**		
Users	**1700**	**1700**		
VTAMs	**1**	**1**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
SFS Servers				
WKSET (V)	3360	3353	-7	-0.21%
TOT CPU/CMD (V)	3.4553	3.4540	-0.0013	-0.04%
CP CPU/CMD (V)	1.5181	1.5268	0.0087	0.57%
VIRT CPU/CMD (V)	1.9372	1.9272	-0.0100	-0.52%
FP REQ/CMD(Q)	1.153	1.156	0.003	0.26%
IO/CMD (Q)	1.638	1.635	-0.003	-0.18%
IO TIME/CMD (Q)	0.022	0.025	0.003	13.64%
SFS TIME/CMD (Q)	0.031	0.037	0.006	19.35%

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

VSE/ESA Guest

This section examines the performance impact of migrating a VSE/ESA guest from VM/ESA 2.3.0 to VM/ESA 2.4.0. All measurements were made on a 9121-320 using the DYNAPACE workload. DYNAPACE is a batch workload that is characterized by heavy I/O. See "VSE Guest (DYNAPACE)" on page 68 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.
3390-2	3990-03	2						1
3390-2	3990-02	4				10	2	
3390-2	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	32MB/XA	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: Performance was equivalent to VM/ESA 2.3.0 for both the V=R and V=V environments.

Table 4. VSE/ESA V=R guest migration from VM/ESA 2.3.0 on the 9121-320

Release Run ID	2.3.0 L1RA8PF2	2.4.0 L1RB8PF2	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 (H)	871.0	841.0	-30.0	-3.44%
ETR (H)	7.72	7.99	0.28	3.57%
ITR (H)	18.83	18.75	-0.08	-0.42%
ITR	18.82	18.58	-0.24	-1.25%
ITRR (H)	1.000	0.996	-0.004	-0.42%
ITRR	1.000	0.988	-0.012	-1.25%
Proc. Usage (Sec)				
PBT/CMD (H)	3.187	3.200	0.013	0.42%
PBT/CMD	3.188	3.229	0.040	1.27%
CP/CMD (H)	0.291	0.280	-0.010	-3.55%
CP/CMD	0.233	0.300	0.067	28.74%
EMUL/CMD (H)	2.896	2.920	0.024	0.82%
EMUL/CMD	2.955	2.928	-0.027	-0.90%
Processor Util.				
TOTAL (H)	40.98	42.62	1.64	4.00%
TOTAL	41.00	43.00	2.00	4.88%
TOTAL EMUL (H)	37.24	38.89	1.65	4.42%
TOTAL EMUL	38.00	39.00	1.00	2.63%
TVR(H)	1.10	1.10	0.00	-0.40%
TVR	1.08	1.10	0.02	2.19%
Storage				
NUCLEUS SIZE (V)	2884KB	2908KB	24KB	0.83%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	38532	38531	-1	0.00%
FREEPGS	85	79	-6	-7.06%
FREE UTIL	0.61	0.54	-0.07	-11.45%
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	7.777	7.509	-0.268	-3.44%
RIO RATE (V)	2.000	2.000	0.000	0.00%
RIO/CMD (V)	15.554	15.018	-0.536	-3.44%
DASD IO TOTAL (V)	350840	348517	-2323	-0.66%
DASD IO RATE (V)	389.82	414.90	25.08	6.43%
DASD IO/CMD (V)	3031.56	3115.46	83.90	2.77%
MDC REAL SIZE (MB)	8.0	7.8	-0.2	-2.93%
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	0.32	0.35	0.03	9.38%
PRIVOPs				
PRIVOP/CMD (R)	10.675	10.473	-0.202	-1.89%
DIAG/CMD (R)	607.677	594.161	-13.517	-2.22%
SIE/CMD	2784.089	2643.143	-140.946	-5.06%
SIE INTCPT/CMD	2310.794	2167.377	-143.417	-6.21%
FREE TOTL/CMD	528.821	503.098	-25.723	-4.86%

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

Table 5 (Page 1 of 2). VSE/ESA V=V guest migration from VM/ESA 2.3.0 on the 9121-320

Release **Run ID**	**2.3.0** **L1VA8PF2**	**2.4.0** **L1VB8PF2**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
Throughput (Min)				
Elapsed Time (H)	483.0	487.0	4.0	0.83%
ETR (H)	13.91	13.80	-0.11	-0.82%
ITR (H)	14.61	14.60	-0.01	-0.04%
ITR	14.65	14.53	-0.12	-0.82%
ITRR (H)	1.000	1.000	0.000	-0.04%
ITRR	1.000	0.992	-0.008	-0.82%
Proc. Usage (Sec)				
PBT/CMD (H)	4.107	4.109	0.002	0.04%
PBT/CMD	4.097	4.131	0.034	0.83%
CP/CMD (H)	1.157	1.173	0.016	1.38%
CP/CMD	1.035	1.087	0.052	5.03%
EMUL/CMD (H)	2.951	2.936	-0.014	-0.49%
EMUL/CMD	3.062	3.044	-0.018	-0.59%
Processor Util.				
TOTAL (H)	95.24	94.50	-0.75	-0.78%
TOTAL	95.00	95.00	0.00	0.00%
TOTAL EMUL (H)	68.42	67.53	-0.89	-1.31%
TOTAL EMUL	71.00	70.00	-1.00	-1.41%
TVR(H)	1.39	1.40	0.01	0.53%
TVR	1.34	1.36	0.02	1.43%
Storage				
NUCLEUS SIZE (V)	2884KB	2908KB	24KB	0.83%
TRACE TABLE (V)	200KB	200KB	0KB	0.00%
PGBLPGS	38459	38452	-7	-0.02%
FREEPGS	107	99	-8	-7.48%
FREE UTIL	0.62	0.59	-0.03	-4.87%
Paging				
PAGE/CMD	112.125	182.625	70.500	62.88%
XSTOR/CMD	0.000	0.000	0.000	na
FAST CLR/CMD	276.000	204.366	-71.634	-25.95%
I/O				
VIO RATE	722.000	718.000	-4.000	-0.55%
VIO/CMD	3113.625	3122.018	8.393	0.27%
RIO RATE (V)	322.000	313.000	-9.000	-2.80%
RIO/CMD (V)	1388.625	1360.991	-27.634	-1.99%
DASD IO TOTAL (V)	153866	149724	-4142	-2.69%
DASD IO RATE (V)	320.55	311.93	-8.63	-2.69%
DASD IO/CMD (V)	1382.39	1356.32	-26.07	-1.89%
MDC REAL SIZE (MB)	112.6	111.7	-0.9	-0.80%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	442	439	-3	-0.68%
MDC WRITES (I/Os)	218	220	2	0.92%
MDC AVOID	407	414	7	1.72%
MDC HIT RATIO	0.85	0.87	0.02	2.35%

Table 5 (Page 2 of 2). VSE/ESA V=V guest migration from VM/ESA 2.3.0 on the 9121-320				
Release **Run ID**	**2.3.0** **L1VA8PF2**	**2.4.0** **L1VB8PF2**	**Difference**	**%Difference**
Environment				
IML Mode	**ESA**	**ESA**		
Real Storage	**160MB**	**160MB**		
Exp. Storage	**0MB**	**0MB**		
VM Mode	**ESA**	**ESA**		
VM Size	**96M**	**96M**		
Guest Setting	**V=V**	**V=V**		
VSE Supervisor	**ESA**	**ESA**		
Processors	**1**	**1**		
PRIVOPs				
PRIVOP/CMD (R)	3117.492	3115.772	-1.720	-0.06%
DIAG/CMD (R)	450.594	444.453	-6.140	-1.36%
SIE/CMD	13618.875	13514.250	-104.625	-0.77%
SIE INTCPT/CMD	11848.421	11892.540	44.119	0.37%
FREE TOTL/CMD	4019.250	3909.045	-110.205	-2.74%
Note: V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

TCP/IP

Telnet, FTP, and NFS measurements were obtained to evaluate the performance effects of migrating from TCP/IP Function Level 310 to TCP/IP Function Level 320. The Telnet and FTP results showed equivalent to slightly improved performance. The NFS results showed significant performance improvements over TCP/IP FL 310. The Telnet, FTP, and NFS results are summarized in the following three sections.

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*	*- Number of Volumes -* *PAGE*	*SPOOL*	*TDSK*	*User*	*Server*	*System*
3390-3	RAMAC 2	4		4		16		
3390-2	3990-2	4	5					
3390-2	3990-3	2	6		1			2 R
3390-2	3990-3	4	5		1			

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: The overall TCP/IP 320 Telnet results are essentially equivalent to the corresponding TCP/IP 310 results. Processor usage in the TCPIP stack virtual machine decreased by 0.5% as a result of the "TCP Layer Pathlength Reductions" item discussed in "Performance Improvements" on page 3.

Table 6 (Page 1 of 2). Telnet migration on a 9121-480

TCP/IP Release **VM/ESA Release** **Run ID**	**310** **2.4.0** **L2BE1802**	**320** **2.4.0** **L2BE1801**	**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	0.126	0.123	-0.003	-2.38%
NONTRIV INT	0.863	0.861	-0.002	-0.23%
TOT INT	0.260	0.257	-0.003	-1.15%
TOT INT ADJ	0.235	0.232	-0.003	-1.39%
AVG FIRST (T)	0.285	0.275	-0.010	-3.34%
AVG LAST (T)	0.398	0.387	-0.010	-2.52%
Throughput				
AVG THINK (T)	24.47	24.53	0.06	0.25%
ETR	57.08	56.96	-0.12	-0.21%
ETR (T)	63.02	63.04	0.02	0.03%
ETR RATIO	0.906	0.904	-0.002	-0.24%
ITR (H)	69.88	69.86	-0.02	-0.03%
ITR	31.68	31.59	-0.09	-0.29%
EMUL ITR	51.81	51.67	-0.14	-0.27%
ITRR (H)	1.000	1.000	0.000	-0.03%
ITRR	1.000	0.997	-0.003	-0.29%
Proc. Usage				
PBT/CMD (H)	28.621	28.629	0.008	0.03%
PBT/CMD	28.561	28.553	-0.008	-0.03%
CP/CMD (H)	11.781	11.791	0.010	0.08%
CP/CMD	11.107	11.104	-0.003	-0.03%
EMUL/CMD (H)	16.839	16.838	-0.002	-0.01%
EMUL/CMD	17.454	17.449	-0.005	-0.03%
Processor Util.				
TOTAL (H)	180.38	180.48	0.10	0.06%
TOTAL	180.00	180.00	0.00	0.00%
UTIL/PROC (H)	90.19	90.24	0.05	0.06%
UTIL/PROC	90.00	90.00	0.00	0.00%
TOTAL EMUL (H)	106.13	106.15	0.02	0.02%
TOTAL EMUL	110.00	110.00	0.00	0.00%
MASTER TOTAL (H)	90.77	90.85	0.08	0.09%
MASTER TOTAL	91.00	91.00	0.00	0.00%
MASTER EMUL (H)	43.96	43.89	-0.07	-0.16%
MASTER EMUL	46.00	46.00	0.00	0.00%
TVR(H)	1.70	1.70	0.00	0.04%
TVR	1.64	1.64	0.00	0.00%
Storage				
NUCLEUS SIZE (V)	2476KB	2476KB	0KB	0.00%
TRACE TABLE (V)	350KB	350KB	0KB	0.00%
WKSET (V)	87	87	0	0.00%
PGBLPGS	55452	55485	33	0.06%
PGBLPGS/USER	30.8	30.8	0.0	0.06%
TOT PAGES/USER (V)	178	178	0	0.00%
FREEPGS	5401	5369	-32	-0.59%
FREE UTIL	0.95	0.95	0.01	0.60%
SHRPGS	1126	1073	-53	-4.71%

Table 6 (Page 2 of 2). Telnet migration on a 9121-480				
TCP/IP Release **VM/ESA Release** **Run ID**	**310** **2.4.0** **L2BE1802**	**320** **2.4.0** **L2BE1801**	**Difference**	**%Difference**
Environment				
Real Storage	**256MB**	**256MB**		
Exp. Storage	**0MB**	**0MB**		
Users	**1800**	**1800**		
VTAMs	**na**	**na**		
VSCSs	**0**	**0**		
Processors	**2**	**2**		
Paging				
READS/SEC	694	698	4	0.58%
WRITES/SEC	422	423	1	0.24%
PAGE/CMD	17.708	17.782	0.074	0.42%
PAGE IO RATE (V)	175.800	174.500	-1.300	-0.74%
PAGE IO/CMD (V)	2.789	2.768	-0.021	-0.77%
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	8.616	8.629	0.013	0.16%
Queues				
DISPATCH LIST	24.11	24.39	0.28	1.16%
ELIGIBLE LIST	0.00	0.00	0.00	na
I/O				
VIO RATE	876	878	2	0.23%
VIO/CMD	13.900	13.927	0.028	0.20%
RIO RATE (V)	619	620	1	0.16%
RIO/CMD (V)	9.822	9.835	0.013	0.13%
NONPAGE RIO/CMD (V)	7.032	7.067	0.034	0.49%
DASD RESP TIME (V)	19.400	19.400	0.000	0.00%
MDC REAL SIZE (MB)	41.4	41.4	-0.1	-0.13%
MDC XSTOR SIZE (MB)	0.0	0.0	0.0	na
MDC READS (I/Os)	175	176	1	0.57%
MDC WRITES (I/Os)	8.55	8.59	0.04	0.47%
MDC AVOID	164	165	1	0.61%
MDC HIT RATIO	0.93	0.93	0.00	0.00%
PRIVOPs				
PRIVOP/CMD	1.633	1.657	0.023	1.43%
DIAG/CMD	36.948	37.118	0.170	0.46%
DIAG 04/CMD	2.529	2.530	0.000	0.02%
DIAG 08/CMD	0.732	0.734	0.002	0.29%
DIAG 0C/CMD	0.192	0.192	0.000	0.21%
DIAG 14/CMD	0.024	0.024	0.000	0.35%
DIAG 58/CMD	1.250	1.250	0.000	0.01%
DIAG 98/CMD	5.179	5.211	0.032	0.61%
DIAG A4/CMD	3.489	3.495	0.006	0.17%
DIAG A8/CMD	2.656	2.663	0.006	0.24%
DIAG 214/CMD	12.687	12.778	0.092	0.72%
DIAG 270/CMD	0.939	0.941	0.002	0.19%
SIE/CMD	65.547	65.719	0.172	0.26%
SIE INTCPT/CMD	44.572	44.689	0.117	0.26%
FREE TOTL/CMD	53.710	53.869	0.159	0.30%
TCPIP Machine				
WKSET (V)	2700	2700	0	0.00%
TOT CPU/CMD (V)	5.4830	5.4550	-0.0280	-0.51%
CP CPU/CMD (V)	2.4420	2.4500	0.0080	0.33%
VIRT CPU/CMD (V)	3.0410	3.0050	-0.0360	-1.18%
DIAG 98/CMD (V)	5.179	5.211	0.032	0.62%
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

FTP

Measurement Description: The measured system consisted of a 9121-480 (2-way) processor running VM/ESA 2.4.0 with TCP/IP Function Level 310 or 320. This system was attached to a 16 Mbit IBM Token Ring using a 3172-3 Interconnect Controller. The TCPIP server was configured to use 32 Kb TCP buffers (DATABUFFERPOOLSIZE) and an MTU size of 4000 bytes. File transfer was to/from an RS/6000 model 250 running AIX 4.2.1 that was connected to the same token ring.

The TCP/IP FL 320 code included APAR PQ28148. This is a change to FTP MODULE which reduces pathlength in the VM client virtual machine for binary get cases.

The workload consisted of a number of consecutive identical FTP file transfers (get or put) initiated by a client virtual machine on the VM/ESA system. For the 2 Mb files, there were 5 file transfers; for the 24 Kb files, there were 100 file transfers. These file transfers were to/from the client virtual machine's 191 minidisk, which was enabled for minidisk caching and defined in RAMAC 2 Array Subsystem storage.

Measurement Results: The measurement results are summarized in Table 7, Table 8, and Table 9. For each table, the absolute results are first shown, followed by the TCP/IP FL 320 to TCP/IP FL 310 ratios.

Table 7. FTP Performance on a 9121-480: Get 2MB Files

Run Id **Mode** **TCP/IP Release**	**G2M_A313** **ASCII** **310**	**G2M_A2F3** **ASCII** **320**	**G2M_B313** **Binary** **310**	**G2M_B2F3** **Binary** **320**
Rate (Kb/sec)	428.5	432.0	454.8	454.8
TCPIP CPU/Kb (msec)	0.274	0.258	0.272	0.260
VM Client CPU/Kb	0.416	0.415	0.079	0.072
Total CPU/Kb	0.690	0.673	0.351	0.332
Rate (Kb/sec)	1.000	1.008	1.000	1.000
TCPIP CPU/Kb (msec)	1.000	0.942	1.000	0.956
VM Client CPU/Kb	1.000	0.998	1.000	0.911
Total CPU/Kb	1.000	0.975	1.000	0.946

Table 8. FTP Performance on a 9121-480: Get 24KB Files

Run Id **Mode** **TCP/IP Release**	**G24KA313** **ASCII** **310**	**G24KA2F3** **ASCII** **320**	**G24KB313** **Binary** **310**	**G24KB2F3** **Binary** **320**
Rate (Kb/sec)	127.2	129.3	150.0	160.9
TCPIP CPU/Kb (msec)	0.885	0.914	0.889	0.893
VM Client CPU/Kb	1.428	1.444	1.617	1.115
Total CPU/Kb	2.313	2.358	2.506	2.008
Rate (Kb/sec)	1.000	1.017	1.000	1.073
TCPIP CPU/Kb (msec)	1.000	1.033	1.000	1.004
VM Client CPU/Kb	1.000	1.011	1.000	0.690
Total CPU/Kb	1.000	1.019	1.000	0.801

Table 9. FTP Performance on a 9121-480: Put 2MB and 24KB Binary Files				
Run Id **TCP/IP Release**	**P2M_B313** **310**	**P2M_B2F3** **320**	**P24KB313** **310**	**P24KB2F7** **320**
Rate (Kb/sec)	496.1	489.1	315.6	315.6
TCPIP CPU/Kb (msec)	0.434	0.440	1.074	1.095
VM Client CPU/Kb	0.064	0.066	1.012	1.029
Total CPU/Kb	0.498	0.506	2.086	2.124
Rate (Kb/sec)	1.000	0.986	1.000	1.000
TCPIP CPU/Kb (msec)	1.000	1.014	1.000	1.020
VM Client CPU/Kb	1.000	1.031	1.000	1.017
Total CPU/Kb	1.000	1.016	1.000	1.018

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: The 6 measured FTP cases showed equivalent or improved performance relative to TCP/IP FL 310. The improvements were confined to the "get" cases. On average for the 4 get cases, throughput improved by 2% while total CPU usage improved by 6%. These improvements resulted from the combined effects of the "TCP Layer Pathlength Reductions" item (applies to gets) and the "FTP Client Pathlength Reduction" item (applies to binary gets). See "Performance Improvements" on page 3 for a discussion of these performance improvements.

NFS

NFS performance has improved significantly as a result of several factors:

- APAR PQ16183, available on TCP/IP Function Level 310, improves the performance of reading and writing Byte File System (BFS) files.
- Additional performance improvements that were incorporated into TCP/IP FL 320 to improve the performance of various BFS and Shared File System (SFS) cases.
- Support for version 3 of the NFS protocol, included in TCP/IP FL 320. This includes support for NFS requests exceeding 8 Kb, piggybacking of file attribute information to minimize the number of trips between client and server, and efficient access of file attributes using the readdirplus function.
- The CMS multitasking improvements included in VM/ESA 2.4.0 reduce NFS server CPU time when accessing BFS files.

This section provides comparison data that quantify these benefits and, additionally, quantify the performance of the TCP support that has been added to NFS in TCP/IP FL 320.

Measurement Description

The measured system consisted of a 9672-R55 (5-way) processor with VM/ESA 2.4.0 running in a shared, 5-way LPAR. This system was attached to a 16 Mbit IBM Token Ring through an OSA-2 adapter. NFS requests were made from an RS/6000 model 250 on the same token ring. The RS/6000 was running AIX 4.2.1, which includes support for both version 2 and version 3 of the NFS protocol. The measured configuration was not dedicated but all measurements were done during periods of negligible additional usage and all measurements were repeated in order to ensure consistency.

All accessed files were on 3390 DASD. Normal recommended CP tuning (large relative share, QUICKDSP ON, minidisk caching enabled, etc) was in effect for the server virtual machines that were active during the measurements (TCPIP, VMNFS, and the SFS/BFS server). The TCPIP server was configured to use an MTU size of 4000 bytes. Except where otherwise stated, the UDP buffer size (LARGEENVELOPEPOOLSIZE) and the TCP buffer size (DATABUFFERPOOLSIZE) were set to 16384.

The following four workloads were used to evaluate NFS performance:

Mount/Umount Mount a 60-file minidisk or directory for binary file access, then unmount it.

Read 200 Kb Do the above mount request, read a 200 Kb file from VM, do the above umount request.

Write 200 Kb Do the above mount request, write a 200 Kb file from VM (replacing the existing file), do the above umount request.

List 60 Files List the names and attributes of the 60 files in the previously mounted minidisk/directory. All files have fixed length records and all file names follow the normal naming conventions for CMS files.

Measurement data were collected using the CP QUERY TIME command, the CP INDICATE USER command, and the QUERY FILEPOOL COUNTER command issued before and after 10 workload invocations. Elapsed times were calculated by the driver program on the AIX client. All results were divided by 10 so that the reported results represent an average of 10 trials.

Mount/umount were included as part of the "Read 200 Kb" workload in order to prevent AIX file caching after the first read invocation from affecting the results. Mount/umount were included in the write workload so as to be consistent with the read workload. The "List 60 Files" workload did not include mount/umount processing as there is no AIX caching in this case.

These 4 workloads were run for files residing on a minidisk, in an SFS directory, and in a BFS directory. The following cases were measured:

- NFS at TCP/IP FL 310, version 2 NFS protocol, UDP transport
- NFS at TCP/IP FL 320, version 2 NFS protocol, UDP transport
- NFS at TCP/IP FL 320, version 3 NFS protocol, UDP transport
- NFS at TCP/IP FL 320, version 3 NFS protocol, TCP transport (minidisk case only)

The desired NFS version and transport protocol were specified on the mount command.

Results: TCP/IP FL 310 NFS V2 to TCP/IP FL 320 NFS V3

Figure 1 and Figure 2 summarize the overall effect of migrating from TCP/IP FL 310 and using the version 2 NFS protocol to TCP/IP FL 320 and using the version 3 NFS protocol for the case of UDP transport. A ratio of 1.00 represents equivalence to TCP/IP FL 310 with V2 NFS. Ratios less than 1.00 represent improved performance. These same results, along with additional measurement data, results for the intermediate case of TCP/IP FL 320 and NFS version 2, and results using the TCP transport protocol are provided in Table 10 on page 35 through Table 17 on page 38.

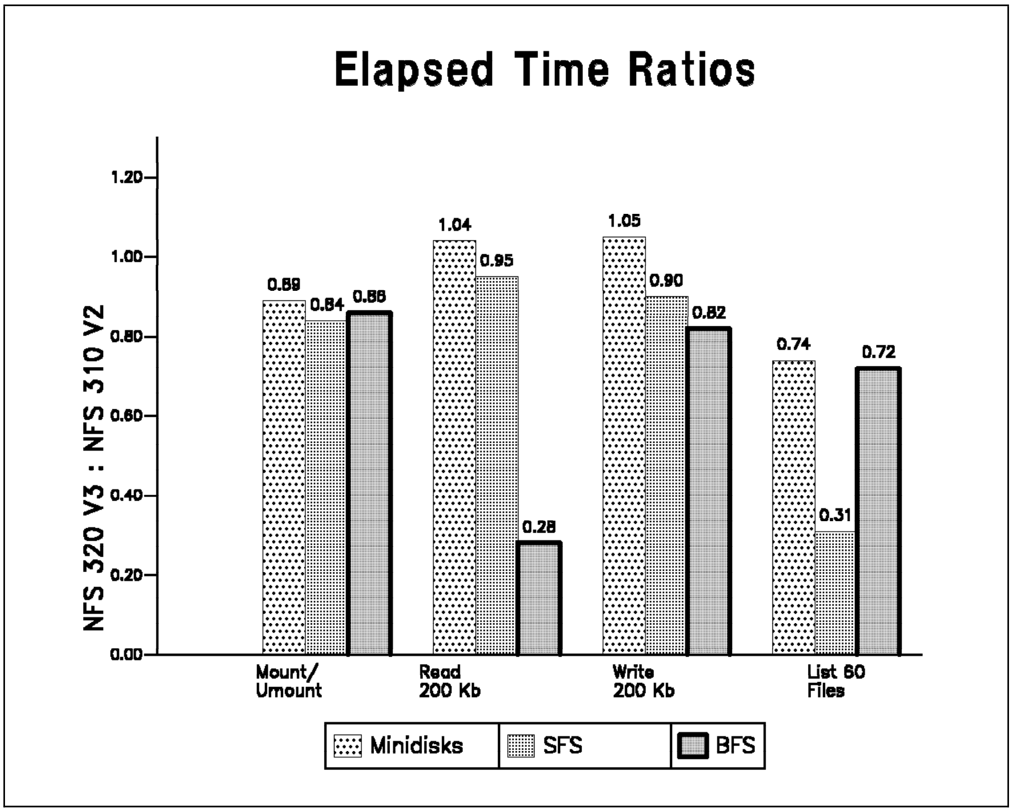

Figure 1. Elapsed times relative to NFS at FL 310 for various NFS workloads

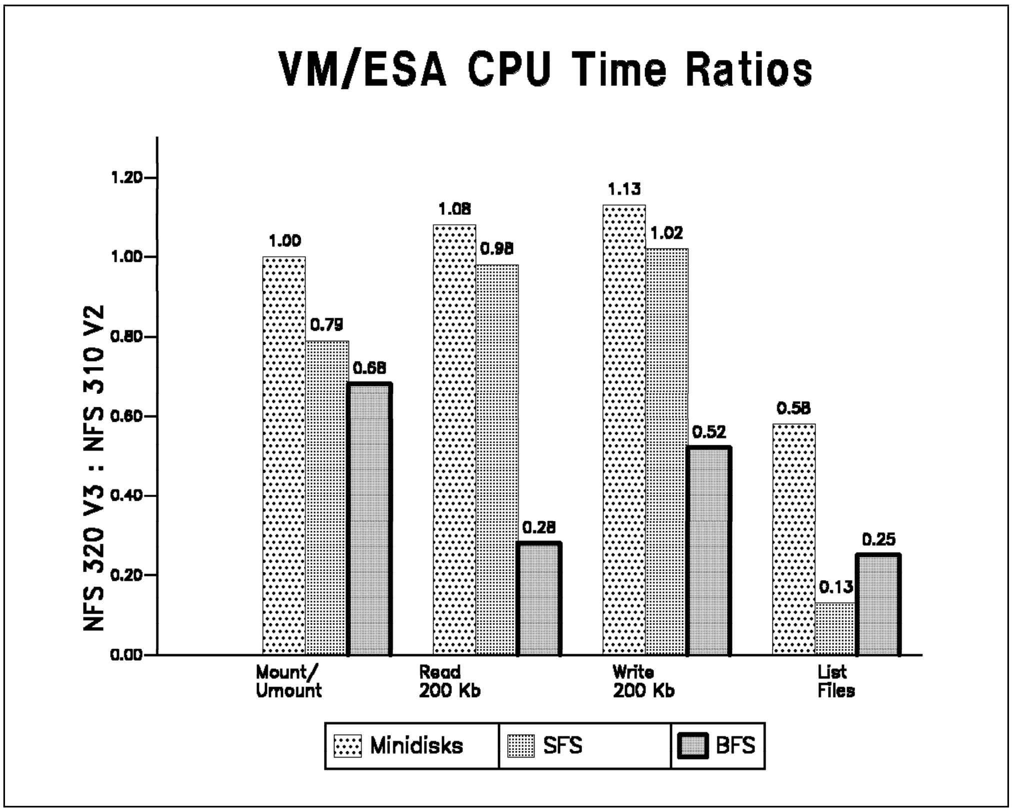

Figure 2. Total VM/ESA CPU times relative to NFS at FL 310 for various NFS workloads

Additional measurements (not shown) confirmed that the performance improvement for the “Mount/Umount” workload is from the mount command. The performance of the umount command has not changed.

The especially large improvement observed for reading and writing BFS files is primarily due to APAR PQ16183, which eliminates large numbers of BFS calls by keeping the file open across NFS requests from the client.

The large improvements observed for the “List 60 Files” workload come from the readdirplus interface that was added to the NFS protocol in version 3. With version 2, the client had to issue a separate NFS request for each object in the mounted directory. With version 3, attribute data for all objects in the directory can be returned to the client in response to a single readdirplus request. For these measurements, the files all had fixed length records. For minidisks or directories containing variable length files, the amount of improvement is much less because the VMNFS server must still open each of the files in order to determine the size of the file.

These results show that most of the performance improvements apply to NFS access to SFS and (especially) BFS files. This has tended to even out NFS performance characteristics for NFS access to files stored in the three different CMS file systems, as illustrated in Figure 3 for elapsed time. A similar uniformity improvement has also occurred for CPU usage.

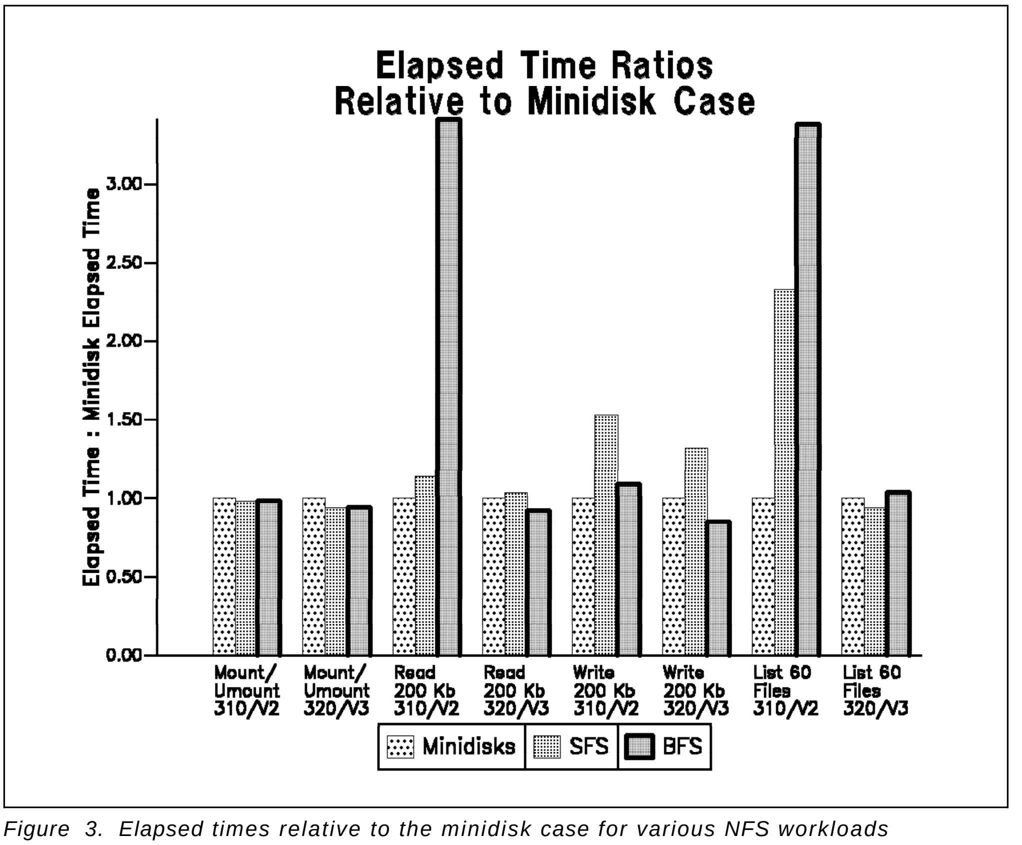

Figure 3. Elapsed times relative to the minidisk case for various NFS workloads

These comparative results do not reflect the ability of the NFS version 3 protocol to support data transfers larger than 8 Kb. These additional benefits occur at larger settings of LARGEENVELOPEPOOLSIZE (UDP) and DATABUFFERPOOLSIZE (TCP) for the TCPIP server virtual machine. See “Results: Large Buffers” on page 39 for additional information.

The CMS multitasking improvements in VM/ESA 2.4.0 (see “Performance Improvements” on page 3) provide additional benefits for NFS access of BFS files by reducing the amount of CPU time used in the VMNFS server virtual machine. These benefits are in addition to the performance benefits reflected in this section, which are all based on measurements done on VM/ESA 2.4.0. For example, comparison of these results to other measurements (not shown) of the 200K BFS read workload running on VM/ESA 2.3.0 CMS reveals an additional 27% decrease in VMNFS server CPU usage and a 14% decrease in total system CPU usage due to the CMS multitasking changes. NFS accesses to minidisk and SFS files are unaffected by these change because the VMNFS server does not use CMS multitasking in those cases.

Measurement results indicate that the combined benefits from support for large data transfers and from the CMS multitasking improvements result in a further 11% elapsed time improvement beyond the 72% elapsed time improvement shown for the case of reading a 200 Kb BFS file in Figure 1, resulting in an overall elapsed time improvement of 83% for this case when all improvements are considered.

Table 10. NFS Performance: Mount/Umount a CMS Minidisk

Run Id **TCP/IP Release** **NFS Protocol Version** **Transport Protocol** **VM File System**	**1M12U161** **310** **V2** **UDP** **Minidisk**	**1M22U162** **320** **V2** **UDP** **Minidisk**	**1M23U162** **320** **V3** **UDP** **Minidisk**	**1M23T162** **320** **V3** **TCP** **Minidisk**
Elapsed Time (sec)	0.41	0.36	0.36	0.37
Elapsed Time Ratio	1.000	0.894	0.891	0.901
CPU Usage				
VMNFS (sec)	0.017	0.019	0.020	0.021
TCPIP (sec)	0.020	0.018	0.017	0.022
Total (sec)	0.037	0.037	0.037	0.043
Total Ratio	1.000	1.000	1.000	1.162
Virtual I/O				
VMNFS	6	5	5	5

Note: 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, binary transfer, 16 Kb transport buffers

Table 11. NFS Performance: Mount/Umount an SFS or BFS Directory

Run Id **TCP/IP Release** **NFS Protocol Version** **Transport Protocol** **VM File System**	**1S12U162** **310** **V2** **UDP** **SFS**	**1S22U162** **320** **V2** **UDP** **SFS**	**1S23U161** **320** **V3** **UDP** **SFS**	**1B12U162** **310** **V2** **UDP** **BFS**	**1B22U161** **320** **V2** **UDP** **BFS**	**1B23U162** **320** **V3** **UDP** **BFS**
Elapsed Time (sec)	0.40	0.35	0.34	0.40	0.36	0.34
Elapsed Time Ratio	1.000	0.875	0.840	1.000	0.894	0.859
CPU Usage						
VMNFS (sec)	0.017	0.017	0.015	0.027	0.022	0.019
TCPIP (sec)	0.022	0.019	0.017	0.021	0.019	0.017
SFS/BFS Server	0.004	0.003	0.002	0.011	0.005	0.004
Total (sec)	0.043	0.039	0.034	0.059	0.046	0.040
Total Ratio	1.000	0.907	0.791	1.000	0.780	0.678
Virtual I/O						
VMNFS	3	3	3	3	3	3
SFS/BFS Server	0	0	0	0	0	0
SFS/BFS Requests	8	7	4	17	7	5
SFS/BFS Time (sec)	0.002	0.002	0.001	0.005	0.004	0.003

Note: 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, binary transfer, 16 Kb transport buffers

“SFS/BFS Requests” and “SFS/BFS Time” are taken from the QUERY FILEPOOL COUNTER data. “SFS/BFS Requests” is from the “Total File Pool Requests” counter, while “SFS/BFS Time” is from the “File Pool Request Service Time” counter and represents elapsed time spent in the SFS/BFS server handling those requests during the measured interval.

Table 12. NFS Performance: Read a 200K Minidisk File				
Run Id	**RM12U162**	**RM22U162**	**RM23U162**	**RM23T162**
TCP/IP Release	**310**	**320**	**320**	**320**
NFS Protocol Version	**V2**	**V2**	**V3**	**V3**
Transport Protocol	**UDP**	**UDP**	**UDP**	**TCP**
VM File System	**Minidisk**	**Minidisk**	**Minidisk**	**Minidisk**
Elapsed Time (sec)	0.74	0.70	0.77	1.03
Elapsed Time Ratio	1.000	0.941	1.040	1.380
CPU Usage				
VMNFS (sec)	0.072	0.078	0.080	0.090
TCPIP (sec)	0.084	0.084	0.089	0.116
Total (sec)	0.156	0.162	0.169	0.206
Total Ratio	1.000	1.038	1.083	1.321
Virtual I/O				
VMNFS	85	84	84	84
Note: 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, binary transfer, 16 Kb transport buffers				

Table 13. NFS Performance: Read a 200K SFS or BFS File						
Run Id	**RS12U162**	**RS22U162**	**RS23U162**	**RB12U162**	**RB22U162**	**RB23U162**
TCP/IP Release	**310**	**320**	**320**	**310**	**320**	**320**
NFS Protocol Version	**V2**	**V2**	**V3**	**V2**	**V2**	**V3**
Transport Protocol	**UDP**	**UDP**	**UDP**	**UDP**	**UDP**	**UDP**
VM File System	**SFS**	**SFS**	**SFS**	**BFS**	**BFS**	**BFS**
Elapsed Time (sec)	0.84	0.80	0.80	2.52	0.69	0.71
Elapsed Time Ratio	1.000	0.952	0.951	1.000	0.274	0.282
CPU Usage						
VMNFS (sec)	0.124	0.116	0.123	0.448	0.119	0.121
TCPIP (sec)	0.062	0.060	0.063	0.068	0.062	0.065
SFS/BFS Server	0.090	0.081	0.084	0.345	0.050	0.052
Total (sec)	0.276	0.257	0.270	0.861	0.231	0.238
Total Ratio	1.000	0.931	0.978	1.000	0.268	0.276
Virtual I/O						
VMNFS	3	3	3	3	3	3
SFS/BFS Server	76	76	76	84	3	2
SFS/BFS Requests	217	190	192	486	66	67
SFS/BFS Time (sec)	0.19	0.20	0.19	1.45	0.08	0.06
Note: 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, binary transfer, 16 Kb transport buffers						

Table 14. NFS Performance: Write a 200K Minidisk File

Run Id **TCP/IP Release** **NFS Protocol Version** **Transport Protocol** **VM File System**	**WM12U162** **310** **V2** **UDP** **Minidisk**	**WM22U162** **320** **V2** **UDP** **Minidisk**	**WM23U162** **320** **V3** **UDP** **Minidisk**	**WM23T162** **320** **V3** **TCP** **Minidisk**
Elapsed Time (sec)	2.79	2.83	2.93	3.00
Elapsed Time Ratio	1.000	1.014	1.050	1.076
CPU Usage				
VMNFS (sec)	0.079	0.091	0.094	0.104
TCPIP (sec)	0.101	0.108	0.109	0.132
Total (sec)	0.180	0.199	0.203	0.236
Total Ratio	1.000	1.106	1.128	1.311
Virtual I/O				
VMNFS	51	50	50	50

Note: 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, binary transfer, 16 Kb transport buffers

Table 15. NFS Performance: Write a 200K SFS or BFS File

Run Id **TCP/IP Release** **NFS Protocol Version** **Transport Protocol** **VM File System**	**WS12U162** **310** **V2** **UDP** **SFS**	**WS22U161** **320** **V2** **UDP** **SFS**	**WS23U162** **320** **V3** **UDP** **SFS**	**WB12U162** **310** **V2** **UDP** **BFS**	**WB22U162** **320** **V2** **UDP** **BFS**	**WB23U162** **320** **V3** **UDP** **BFS**
Elapsed Time (sec)	4.28	3.76	3.86	3.03	2.68	2.48
Elapsed Time Ratio	1.000	0.879	0.901	1.000	0.884	0.818
CPU Usage						
VMNFS (sec)	0.237	0.239	0.241	0.564	0.266	0.280
TCPIP (sec)	0.070	0.071	0.075	0.073	0.072	0.076
SFS/BFS Server	0.251	0.246	0.252	0.459	0.200	0.212
Total (sec)	0.558	0.556	0.568	1.096	0.538	0.568
Total Ratio	1.000	0.996	1.018	1.000	0.491	0.518
Virtual I/O						
VMNFS	3	3	3	3	3	3
SFS/BFS Server	260	257	261	83	83	81
SFS/BFS Requests	485	452	456	623	201	204
SFS/BFS Time (sec)	3.50	3.04	3.14	1.78	1.86	1.65

Note: 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, binary transfer, 16 Kb transport buffers

Table 16. NFS Performance: List Contents of a 60-File Minidisk

Run Id	**LM12U162**	**LM22U162**	**LM23U162**	**LM23T162**
TCP/IP Release	**310**	**320**	**320**	**320**
NFS Protocol Version	**V2**	**V2**	**V3**	**V3**
Transport Protocol	**UDP**	**UDP**	**UDP**	**TCP**
VM File System	**Minidisk**	**Minidisk**	**Minidisk**	**Minidisk**
Elapsed Time (sec)	0.24	0.23	0.18	0.22
Elapsed Time Ratio	1.000	0.970	0.738	0.907
CPU Usage				
VMNFS (sec)	0.009	0.009	0.007	0.007
TCPIP (sec)	0.015	0.015	0.007	0.008
Total (sec)	0.024	0.024	0.014	0.015
Total Ratio	1.000	1.000	0.583	0.625
Virtual I/O				
VMNFS	6	6	6	6

Note: 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, binary transfer, 16 Kb transport buffers

Table 17. NFS Performance: List Contents of a 60-File SFS or BFS Directory

Run Id	**LS12U161**	**LS22U162**	**LS23U162**	**LB12U162**	**LB22U162**	**LB23U162**
TCP/IP Release	**310**	**320**	**320**	**310**	**320**	**320**
NFS Protocol Version	**V2**	**V2**	**V3**	**V2**	**V2**	**V3**
Transport Protocol	**UDP**	**UDP**	**UDP**	**UDP**	**UDP**	**UDP**
VM File System	**SFS**	**SFS**	**SFS**	**BFS**	**BFS**	**BFS**
Elapsed Time (sec)	0.56	0.63	0.17	0.81	0.64	0.58
Elapsed Time Ratio	1.000	1.133	0.306	1.000	0.789	0.719
CPU Usage						
VMNFS (sec)	0.070	0.086	0.009	0.166	0.048	0.039
TCPIP (sec)	0.052	0.065	0.008	0.031	0.025	0.017
SFS/BFS Server	0.042	0.052	0.005	0.121	0.026	0.023
Total (sec)	0.164	0.203	0.022	0.318	0.099	0.079
Total Ratio	1.000	1.238	0.134	1.000	0.311	0.248
Virtual I/O						
VMNFS	0	0	0	0	0	0
SFS/BFS Server	2	2	1	0	1	1
SFS/BFS Requests	103	121	9	198	36	33
SFS/BFS Time (sec)	0.04	0.04	0.01	0.06	0.03	0.03

Note: 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, binary transfer, 16 Kb transport buffers

Results: TCP Support

NFS at TCP/IP FL 310 only supported the UDP protocol. With TCP/IP FL 320, NFS has added support for the TCP protocol. A set of measurements was obtained to compare NFS performance when using the UDP and TCP transport protocols. These measurements were obtained for the case of binary access to minidisk files using TCP/IP FL 320 and NFS version 3. The results are summarized below, with further details provided in Table 10, Table 12, Table 14, and Table 16.

Table 18. Effect of transport protocol on NFS performance

Measure **Transport Protocol**	**ET** **UDP**	**ET** **TCP**	**ET** **Pct**	**TCPU** **UDP**	**TCPU** **TCP**	**TCPU** **Pct**
Mount/Umount	0.36	0.37	3	0.037	0.043	16
Read 200 Kb	0.77	1.03	34	0.169	0.206	22
Write 200 Kb	2.93	3.00	2	0.203	0.236	16
List 60 Files	0.18	0.22	22	0.014	0.015	7

Note: ET = Elapsed Time (sec), TCPU = Total system CPU usage (sec); 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, TCP/IP 320, NFS version 3, minidisk files, binary transfer, 16 Kb transport buffers

These results indicate that NFS performance tends to be somewhat better when the UDP protocol is used.

These results are for the case of NFS access to minidisk files. The UDP to TCP performance relationship should be similar for NFS access to SFS or BFS files.

Results: Large Buffers

When reading or writing large files, NFS performance is significantly affected by the number of bytes requested by the client per request. The larger this is, the fewer NFS requests are required, and the better the performance.

When the NFS version 2 protocol is used by an NFS client, the VMNFS server can read or write a maximum of 8 Kb. This is a limitation of the version 2 protocol and applies to both the UDP and TCP transport protocols. With version 3, this 8 Kb limit has been removed, creating the potential for improved performance when reading or writing large files.

For the NFS version 3 protocol, the VMNFS server defines maximum and preferred read and write data transfer sizes based on TCP/IP transport layer buffer sizes defined in the PROFILE TCPIP (or nodename TCPIP) configuration file. For UDP-connected clients, the maximum and preferred read and write sizes are based on the buffer size specified on the LARGEENVELOPEPOOLSIZE statement. For TCP-connected clients, these are based on the buffer size specified on the DATABUFFERPOOLSIZE statement[2].

When using the NFS version 3 protocol, NFS clients can determine these maximum and preferred sizes and use this information to help determine the actual data transfer sizes they should use when making requests to the VM NFS server. Exactly how this information is used will vary from one client implementation to another.

[2] For DATABUFFERPOOLSIZE values greater than 64 Kb, the VMNFS server uses 64 Kb to compute the maximum and preferred read and write sizes. The largest DATABUFFERPOOLSIZE value supported by TCP/IP for VM is 256 Kb.

For additional information, refer to the "Configuring the NFS Server" chapter in the *TCP/IP Function Level 320 Planning and Customization* manual.

Two sets of measurements were obtained for NFS binary read of a 200 Kb minidisk file using various TCP/IP transport layer buffer sizes. For one set, the NFS client used the UDP transport protocol, while for the other set, TCP was used. The results are summarized in the following table.

Table 19. NFS read of a 200Kb minidisk file: Effect of UDP/TCP buffer size

Run Id **Transport Protocol** **Buffer Size (bytes)**	**RM23U082** **UDP** **8192**	**RM23U164** **UDP** **16384**	**RM23U321** **UDP** **32768**	**RM23T082** **TCP** **8192**	**RM23T164** **TCP** **16384**	**RM23T323** **TCP** **32768**
Maximum/Preferred (bytes)	8784	15952	32336	7784	15976	32360
Max Transfer (bytes)	8320	8320	16512	4224	8320	16512
Elapsed Time (sec)	0.76	0.76	0.71	1.23	1.01	1.05
Elapsed Time Ratio	1.000	0.990	0.925	1.000	0.818	0.849
CPU Usage						
VMNFS (sec)	0.073	0.071	0.056	0.115	0.083	0.065
TCPIP (sec)	0.079	0.078	0.058	0.163	0.109	0.079
Total (sec)	0.152	0.149	0.114	0.278	0.192	0.144
Total Ratio	1.000	0.980	0.750	1.000	0.691	0.518
Virtual I/O						
VMNFS	84	84	84	84	84	84
TCP/IP	931	925	682	2646	1581	1123
TCP/IP I/O Ratio	1.000	0.994	0.733	1.000	0.598	0.424

Note: 9672-R55 shared LPAR, VM/ESA 2.4.0, 16 Mbit IBM Token Ring, OSA-2 Adapter, AIX 4.2.1 client on RS/6000 250, binary transfer TCP/IP 320, NFS version 3

"Maximum/Preferred" refers to the maximum and preferred values that were selected by the VMNFS server. This information was obtained by issuing the SMSG VMNFS M Q CONFIG command. The maximum read, maximum write, and preferred values were the same.

The maximum number of bytes transferred per NFS read request ("Max Transfer") was determined for each case. This was done by looking at the lengths returned to the client in a 501 trace (SMSG VMNFS M MASK 501). These lengths include both data and header bytes.

Note that the maximum and preferred values selected by the VMNFS server are normally slightly smaller than the corresponding TCP/IP stack configuration values. This is to allow space for the RPC header information that is required on all data transfers. The UDP 8192 case is an apparent exception to this. When a LARGEENVELOPEPOOLSIZE of 8192 bytes is specified, the TCP/IP stack actually creates 9216-byte (9 Kb) UDP buffers in order to accommodate the special requirements of the NFS version 2 protocol and the ATM protocol. It is this 9216 value that the TCP/IP stack communicates to the VMNFS server, resulting in the observed 8784 maximum and preferred values.

Note that, except for the UDP 8192 case, the maximum number of bytes transferred per NFS read request is only about half the size of the maximum/preferred values advertised by the VMNFS server. This result illustrates the fact that the actual number of bytes per read is determined by the client. When our AIX client sees that the maximum/preferred value is not, for example, a full 16 Kb, it decides to receive data in increments of half that size. Other clients may instead, for example, choose to use the maximum/preferred value or drop down to the next 1 Kb multiple.

When the transport buffer is specified as 8 Kb, NFS with UDP performs especially well relative to NFS with TCP. This is because NFS with UDP gets to transfer data 8 Kb at a time (due to the special case 9 Kb transport buffer mentioned above), while NFS with TCP transfers data 4 Kb at a time because the client drops down to half the buffer size when it finds that the full 8 Kb is not available.

As Max Transfer increases, the number of NFS requests needed to transfer the data decreases. This causes TCPIP I/Os, TCPIP CPU usage, and VMNFS CPU usage to also decrease. This improved performance is possible because the NFS version 3 protocol supports data transfer in increments larger than 8 Kb.

Migration from Other VM Releases

The performance results provided in this report apply to migration from VM/ESA 2.3.0. 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*

240 *VM/ESA Version 2 Release 4.0 Performance Report* (this document)

See “Referenced Publications” on page vii 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 2.2.0 to VM/ESA 2.4.0 will tend to affect VSE guest performance, look at the VM/ESA 2.2.0 to VM/ESA 2.3.0 comparison measurements and the VM/ESA 2.3.0 to VM/ESA 2.4.0 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 44.

Table 20. 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	230			
		9121-480	230			
		9121-320			230	
VM/ESA 2.3.0	VM/ESA 2.4.0	9121-742	240			
		9121-480	240			
		9121-320			240	

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 21 summarizes the approximate ITR relationships for the CMS-intensive environment for migrations to VM/ESA 2.4.0.

Table 21. Approximate VM/ESA 2.4.0 relative capacity: CMS-intensive environment

Source	Case	ITRR	Notes
VM/SP 5	9221-120	0.90	1,2,5-7
VM/SP 6	9221-120	1.05	2,5-7
VM/ESA 1.0 (370)	9221-120 9221-170	0.98 1.05	2,5-7 4-7
VM/ESA 1.5 (370)	9221-120 9221-170	0.96 1.03	2,5-7 4-7
VM/SP HPO 5	UP MP	0.99 1.10	4,6,7 3,4,6,7
VM/XA 2.0		1.22	7
VM/XA 2.1		1.19	7
VM/ESA 1.0 ESA		1.15	7
VM/ESA 1.1		1.10	7
VM/ESA 2		1.09	7
VM/ESA 2.1		1.08	7
VM/ESA 2.2		1.05	7
VM/ESA 2.1.0		1.00	
VM/ESA 2.2.0		0.99	
VM/ESA 2.3.0		0.99	

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 2.4.0 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 2.4.0 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 2.4.0 system has additional real storage.

   For example, the stated VM/SP HPO 5 to VM/ESA 2.4.0 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 21 on page 44 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 2.4.0. From Table 21 on page 44, the estimated ITRR for migrating from VM/ESA 2.1 to VM/ESA 2.4.0 is 1.08. Therefore, the estimated overall increase in system capacity is 1.30*1.08 = 1.40.

Table 21 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.

Additional Evaluations

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

LIMITHARD Share Capping Option

The LIMITHARD option of the SET SHARE command has been available since VM/ESA 1.2.2. Use of the LIMITHARD option allows a cap or limit to be placed the processor resources used by a virtual machine. While LIMITHARD proved to be very effective for most customers, it was less accurate in certain environments. In particular, the LIMITHARD capping seemed overly aggressive for systems where the total processor utilization was low. In VM/ESA 2.4.0, changes have been made to improve the tracking of the limit share values for the LIMITHARD option. Figure 4 illustrates the improvements.

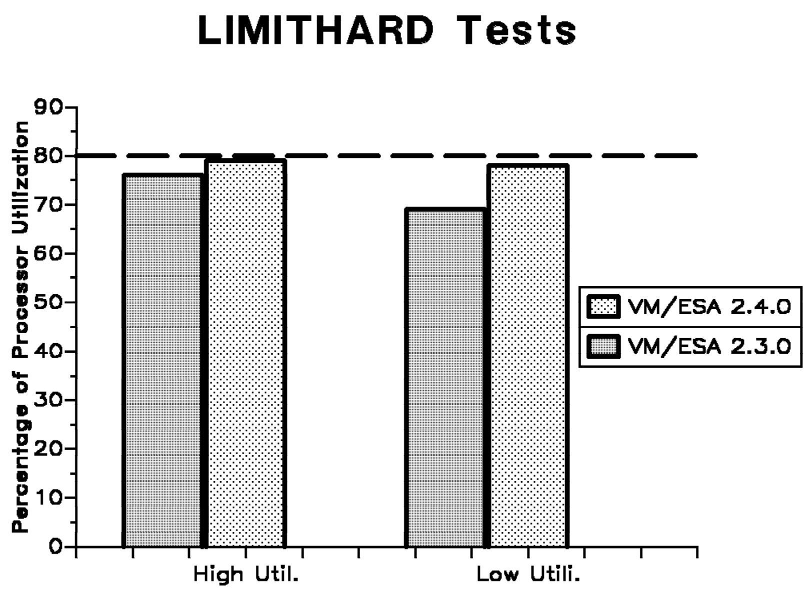

Figure 4. LIMITHARD tests with a user share setting of REL 100 ABS 20% LIMITHARD on a 9121-742.

In the tests shown above, a processor intensive job was run in a user that had an absolute 20% hard limit. The absolute setting is of the system, which in this case had four processors (9121-742). Therefore the target maximum was 80% of a single processor. The dashed line in Figure 4 represents this target maximum. Two scenarios were run for each release. In the first one, the overall system utilization was above 95%, while in the second one the overall system utilization was below 30%. For the VM/ESA 2.3.0 runs, the user is limited more than the hard limit, while the VM/ESA 2.4.0 user is held very close to the target maximum.

The limit hard accuracy has been improved. However, in an LPAR environment, results can still be unpredictable. Tests performed at IBM showed that in a shared LPAR environment the user was limited more the the limit hard setting. One early support program customer saw the opposite results. In their LPAR environment, the limited user received more than the target maximum.

A side effect of the improved share limit accuracy is that the system overhead increases when a limited user is being held back. The amount of overhead is also dependent on the minor dispatch slice. For the tests run above with a default minor dispatch slice, system processor usage increased from 0.4% to 1.9%.

Reusable Server Kernel

The Reusable Server Kernel (RSK) is an execution framework that facilitates the writing of high performance servers on VM/ESA. The RSK provides functionality that is needed by all or many VM servers, thus freeing the server writer to concentrate on providing the application-specific functions. This section takes a look at the performance of some of the RSK functions.

The RSK code, first made available on VM/ESA 2.2.0 and VM/ESA 2.3.0 via APAR VM61878, is now shipped with VM/ESA 2.4.0. The user′s manual, sample server, and supplementary information are available at http://www.ibm.com/s390/vm/rsk.

The measurements make use of the ECHO and SGEXER commands provided in the SSTEST sample server. They also made use of some thin scaffolding built around the ssWorkerAllocate API.

All measurements were obtained on a 9121-480 (2-way) processor running VM/ESA 2.3.0. Each measurement consisted of 1 to 8 equivalent client virtual machines sending commands to the RSK server virtual machine using the CP MSG command. After all clients were started, this activity was measured for a fixed time interval. Commands executed and average response time were collected for each client and aggregated over all clients. In addition, CP monitor records were collected for each measurement and reduced by VMPRF.

Echo Results

The SSTEST ECHO command simply responds to the client with the character string that was sent. This command provides a good measure of baseline performance (no application content) and exercises RSK code that is common to all RSK server requests. Table 22 contains results for this workload at various numbers of clients. For these measurements, each client waited 5 msec between echo commands, the measurement interval was 1 minute, and 3 bytes were sent/returned.

Table 22. RSK Performance: Echo Command

clients **run id**	**1** **E005C1A**	**4** **E005C4A**	**6** **E005C6A**	**8** **E005C8A**
commands/sec	56.0	163.9	190.8	192.7
response time (msec)	4.5	9.1	13.8	24.4
total CPU utilization	27.7	77.3	89.5	91.8
RSK server				
Total CPU utilization	18	52	62	64
Total CPU/cmd (msec)	3.2	3.2	3.2	3.3
Virtual CPU/cmd (msec)	2.1	2.1	2.2	2.3
CP CPU/cmd (msec)	1.1	1.1	1.0	1.0

Note: 9121-480, VM/ESA 2.3.0, 5 msec think time

The results indicate that:

- Throughput is limited by processor usage and waiting for processor time. At 8 clients, total processor utilization was about 92%. Processor utilization peaked at significantly less than 100% because CP MSG was used for client/server communications. Some of the CP code that implements the CP MSG facility must run on the master processor. As a result, the client and server virtual machines spent some of their time waiting to get the master processor. This shows up as ″waiting on console function″ in the USER_STATES VMPRF report. If, for example, the IUCV line driver were used instead, maximum processor utilization should be nearly 100% because none of the IUCV code requires the master processor.
- Processor usage per command is essentially independent of load.

These results were compared to similar results collected for echo requests sent to an SFS server virtual machine, as approximated by the QUERY FILEPOOL CONNECT command. Those results were similar except that 1) maximum CPU utilization was over 99% because APPC was used for communications (like IUCV, APPC does not require the master processor) and 2) server total CPU time per command was about 63% lower. RSK uses the generalized CMS multitasking interfaces to achieve multithreading, while SFS uses its own subdispatching mechanism that has been tailored to its requirements.

Storage Group Results

The SSTEST SGEXER command was used to write and then read the first 50 storage group blocks, 10 blocks at a time. Measurements were made with 1 and 2 clients. For the 2-client measurements, each client used a different storage group and those storage groups were implemented on separate DASD volumes. This was done in order to verify that an I/O by one client does not block the other client from continuing to run and doing an I/O at the same time. A pair of measurements was obtained for the following two cases:

- The storage groups were implemented as VM data spaces mapped to minidisks that were disabled for minidisk caching.
- The storage groups were implemented as minidisks enabled for minidisk caching and accessed using asynchronous Diagnose X′250′.

The measurement results are summarized in Table 23.

Table 23. RSK Performance: Storage Group Usage

storage group type **minidisk caching** **clients** **run id**	**DS** **disabled** **1** **GTDS1M0A**	**DS** **disabled** **2** **GTDS2M0A**	**D250** **enabled** **1** **GTBS1A**	**D250** **enabled** **2** **GTBS2B**
commands/sec	8.63	16.78	8.42	16.28
response time (msec)	115.4	118.8	118.3	122.4
total CPU utilization	16.0	27.9	17.5	30.5
RSK server				
Total CPU/cmd (msec)	23.2	22.6	26.1	27.0
Virtual CPU/cmd (msec)	16.2	16.7	14.3	13.5
CP CPU/cmd (msec)	7.0	6.0	11.9	13.5
DASD IOs/sec	43.4	84.7	42.3	82.0
DASD IOs/cmd	5.0	5.0	5.0	5.0
DASD response (msec)	17.8	17.9	17.8	17.9

Note: 9121-480, VM/ESA 2.3.0, no think time, write and then read 50 blocks, 10 blocks per I/O, each client uses a separate storage group on a separate DASD volume, no paging

The fact that throughput doubles when going from 1 to 2 clients indicates that, in both the data space and Diag 250 cases, there is effective overlap of processing for the two clients. It also indicates that throughput would have gone substantially higher had additional, equivalently configured clients and storage groups been present.

The data space case was run with minidisk caching disabled. This is the recommended tuning for minidisk-mapped data spaces because it avoids the extra overhead of caching the file data twice.

Performance of the data space and the Diag 250 cases was similar, with somewhat lower processor usage for the data space case. Data spaces are a good choice when the number of frequently accessed blocks is sufficiently small relative to real memory size that little paging occurs. When a page I/O occurs, all threads in the server machine are blocked until that I/O completes. Therefore, Diag 250 is a better choice when large amounts of data are involved.

Worker Machine Results

To evaluate the RSK′s worker machine function, some scaffolding was build around ssWorkerAllocate such that a CP MSG to the RSK would cause a very brief worker transaction. ssWorkerAllocate would be called to connect to a worker, the worker would be sent 4 Kb of data (which it would discard), and then the connection to the worker would be closed. Resource consumption by this brief transaction is indicative of the RSK′s overhead in selecting a worker machine and establishing a connection to it.

Table 24 summarizes the results obtained for the case of 1 to 3 clients exercising this workload with a 5 second think time between requests. The 5 available worker machines were all logged on before the measurement.

Table 24. RSK Performance: Echo to Worker Machines

clients **run id**	**1** **W05C1A**	**2** **W05C2A**	**3** **W05C3A**
commands/sec	34.3	35.0	45.1
response time (msec)	15.0	42.0	48.5
total CPU utilization	65.6	66.7	89.7
RSK server			
Total CPU utilization	46	46	60
Total CPU/cmd (msec)	13.4	13.3	13.4
Virtual CPU/cmd (msec)	9.3	9.3	9.4
CP CPU/cmd (msec)	4.0	4.0	4.0
workers			
Total CPU/cmd (msec)	15.1	15.1	17.4
Virtual CPU/cmd (msec)	8.0	8.1	9.0
CP CPU/cmd (msec)	7.0	7.0	8.3
total			
Total CPU/cmd (msec)	28.5	28.4	30.7
Virtual CPU/cmd (msec)	17.4	17.3	18.4
CP CPU/cmd (msec)	11.1	11.0	12.3

Note: 9121-480, VM/ESA 2.3.0, 5 second think time, 5 worker machines

The achievement of higher throughputs when the number of clients is increased verifies that the client virtual machines, the RSK server, and the worker virtual machines are able to do work concurrently with a high level of independence.

These results and the USER_STATES VMPRF report suggest that maximum throughput is being limited by the capacity of the RSK server which, in turn, is limited by processor usage, waiting for processor time, and waiting for the master processor. Total processor utilization of nearly 90% for the 3-client case shows that, additionally, throughput is close to being limited by total processing capacity.

These results show higher processor usage per command than the echo results presented in Table 22 as a result of using a separate worker virtual machine to perform the echo function.

Tivoli ADSM Version 3

Tivoli ADSM for VM Version 3 includes a number of significant performance improvements. This section provides measurement results that assess the performance differences between ADSM Version 2 and Version 3. Tuning information is also provided.

Measurement Description

The ADSM server virtual machine (DSMSERV) was run on a 9121-480 (2-way) processor running VM/ESA 2.3.0 and TCP/IP Function Level 310. The files to be backed up or restored were on one or two RS/6000 model 250 client workstations. The workstation used for the single client runs was running AIX 4.2.1, while the workstation that was added for the 2-client runs was running AIX 4.1.4. These workstations were connected to the VM/ESA system via 16 Mbit IBM Token Ring. The VM/ESA system was connected to the token ring through a 3172-3 communications controller.

Measurements were made using ADSM for VM Version 2 (V2) and Version 3 (V3). For all measurements, the version of ADSM client code used was the same as the version of the ADSM server code being run on VM/ESA.

Three different workloads were measured — backup of files from 1 client workstation to a disk storage pool, backup of files from 2 client workstations to a disk storage pool, and restore of those files for 1 client to that workstation. For each of these cases, a separate measurement was obtained for each of 5 different file sizes — 1KB, 10KB, 100KB, 10Mb, and 256Mb. This is the same workload that is used by the ADSM performance team to evaluate ADSM performance on other platforms. No earlier backup copies of the files being backed up were already in the storage group. Likewise, the files being restored did not already exist on the client.

CP monitor data (later reduced by VMPRF) and the AIX console were collected for each measurement. MB transferred and elapsed time were taken from the AIX console and used to calculate KB/sec. The remaining data were obtained from the VMPRF reports or derived from a combination of the VMPRF data and the AIX console information.

Preliminary Tuning Experiments

These workloads were first run with no special tuning in either the client or server systems. Several experiments were then performed using backup of a 10M file as the workload in order to assess the impact of various tuning changes. This resulted in a more optimal set of tuning settings, which were then used to obtain the results provided here.

This tuning is summarized in Table 25, which shows the default values, the tuning settings used for the "untuned" baseline runs (often the default values), the final tuning settings used for the tuned runs, and the approximate degree of benefit relative to the untuned setting that was observed for the test workload.

Table 25. Tuning settings used for untuned and tuned ADSM measurements

Tuning Parameter	default	untuned	tuned	benefit
AIX client: ADSM				
tcpbufsize	8	8	32	some
tcpwindowsize	16	16	64	large
txnbytelimit	200	200	25600	?
AIX client: TCP/IP				
tcp_sendspace	4096	16384	65536	large
tcp_recvspace	4096	16384	65536	large
MTU size	1492	1492	2000	some
VM: TCPIP server				
databufferpoolsize	8192	8192	65536	large
MTU size	576	4000	4000	n/c
VM: ADSM server				
txngroupmax	16	16	256	?

A "?" in the benefit column indicates a tuning change that was made based on a documented recommendation, but not verified by these informal measurements.

The client ADSM options are in the dsm.sys file. The AIX client TCP/IP parameters are set using the "no" AIX command, except for MTU, which is set using "smit tcpip." The TCP/IP VM parameters are set in the PROFILE TCPIP file, while the ADSM for VM server options are in the DSMSERV OPT file.

In addition to the tuning listed above, the normal CP tuning recommended for server virtual machines was done for the TCPIP and DSMSERV server virtual machines. This included the QUICKDSP option, larger than default (100) relative SHARE settings, and use of minidisk caching for DSMSERV's minidisks.

The combined effect of these tuning actions on overall performance was substantial. Throughput increased by 17% to 190%, while total VM/ESA CPU usage decreased by 5% to 40% relative to the untuned configuration. The larger the file size, the larger the improvement.

It is interesting to note that even though the tuned configuration showed much greater absolute performance than the untuned configuration, both configurations showed a very similar pattern and degree of performance improvement when going from ADSM V2 to V3. This indicates that the V3 improvements are likely to express themselves to a similar degree across a wide variety of ADSM configurations.

Results Summary

Performance was nearly always better with ADSM V3, both in terms of throughput and VM CPU and I/O resource usage. For backup, throughput improvements ranging from 7% to 188% were observed, while VM CPU and I/O usage [3] improved by 17% to 66% and by 77% to 87% respectively. Restore improvements were smaller. Throughput improvements ranged from -3% to

15%, while VM CPU and I/O usage improved by 21% to 52% and by 0% to 47% respectively.

For both backups and restores, the largest percentage improvements were observed for the smallest (1KB) files. For small files, the performance improvements are mostly due to a V3 enhancement known as "server file aggregation" which groups together data from multiple small files and transfers them as a single unit. Performance also improved for large files, although to a smaller extent, primarily due to the new "uselargebuffers" server option (which is the default).

Environments characterized by high levels of DSMSERV loading that result in processor and/or I/O contention will tend to experience higher levels of improvement when migrating from V2 to V3 than those shown in this report. This higher degree of improvement will result from the large decreases in CPU and I/O usage that were achieved in V3. See "Backup Results: 2 Clients" on page 55 for further discussion.

In nearly all cases, throughput was primarily limited by the relatively high latencies associated with the network hardware configuration that was used for these measurements. Faster network configurations can be expected to deliver higher throughputs than those observed here.

The tuned configuration results for 1 client backup, 2 client backup, and 1 client restore are presented and discussed in the following three sections.

---

[3] CPU usage refers to total system CPU time per KB transferred. I/O usage refers to the count of virtual DASD I/Os issued by the DSMSERV server virtual machine per KB transferred. Because of minidisk caching, not all of these virtual I/Os result in real I/Os.

Backup Results: 1 Client

Table 26. Backup 1 client: ADSM V2

File Size	1KB	10KB	100KB	10MB	256MB
MB Transferred	12.2	20.0	50.0	100.0	257.9
Elapsed Time (sec)	410	83	86	147	371
CPU Utilization	37.2	36.3	17.5	14.4	13.3
DASD I/O rate (/sec)	34.7	30.7	28.1	23.8	24.8
DASD Response (msec)	8.1	12.0	13.2	14.8	14.9
DSMSERV CPU Util	64	59	16	10	7
TCPIP CPU Util	8	11	16	16	17
KB/sec	31	247	597	697	712
Total CPU/KB (msec)	24.39	2.94	0.59	0.41	0.37
DASD IOs/KB	1.14	0.12	0.05	0.03	0.03

Note: 9121-480, VM/ESA 2.3.0, 16 Mbit IBM Token Ring, 3172-3, AIX 4.2.1 client on RS/6000 250

Table 27. Backup 1 client: ADSM V3

File Size	1KB	10KB	100KB	10MB	256MB
MB Transferred	12.3	20.0	50.0	100.0	258.0
Elapsed Time (sec)	182	52	80	134	341
CPU Utilization	31.4	23.6	13.9	11.7	11.9
DASD I/O rate (/sec)	10.8	9.0	6.4	5.6	5.5
DASD Response (msec)	19.5	37.1	56.3	74.0	75.4
DSMSERV CPU Util	42	28	9	3	4
TCPIP CPU Util	18	18	16	18	18
KB/sec	69	397	639	767	775
Total CPU/KB (msec)	9.11	1.19	0.43	0.31	0.31
DASD IOs/KB	0.16	0.02	0.01	0.01	0.01

Note: 9121-480, VM/ESA 2.3.0, 16 Mbit IBM Token Ring, 3172-3, AIX 4.2.1 client on RS/6000 250

Table 28. Backup 1 client: V3/V2 ratios

File Size	1KB	10KB	100KB	10MB	256MB
MB Transferred	1.00	1.00	1.00	1.00	1.00
Elapsed Time (sec)	0.44	0.62	0.93	0.91	0.92
CPU Utilization	0.84	0.65	0.79	0.81	0.89
DASD I/O rate (/sec)	0.31	0.29	0.23	0.24	0.22
DASD Response (msec)	2.41	3.09	4.27	5.00	5.06
DSMSERV CPU Util	0.66	0.47	0.54	0.33	0.49
TCPIP CPU Util	2.26	1.56	1.00	1.13	1.06
KB/sec	2.26	1.60	1.07	1.10	1.09
Total CPU/KB (msec)	0.37	0.41	0.74	0.74	0.82
DASD IOs/KB	0.14	0.18	0.21	0.21	0.20

Note: 9121-480, VM/ESA 2.3.0, 16 Mbit IBM Token Ring, 3172-3, AIX 4.2.1 client on RS/6000 250

The CPU and DASD I/O data in these tables and in the tables to follow pertain to the VM/ESA system. Nearly all of the CPU usage is from the DSMSERV and TCPIP virtual machines, while nearly all of the DASD I/Os are from the DSMSERV virtual machine.

Note that ADSM throughput and efficiency improves greatly with increasing file size. There is a significant amount of processing required for each file that is independent of file size. This constant amount of per-file processing becomes an insignificant part of total processing as file size becomes very large. The “server file aggregation” improvement reduces this per-file processing and, as such, is a key improvement for small files but has no effect on the performance of large file transfers.

Note that DASD response time increases as DASD I/Os per KB decrease since more data, on average, is being transferred per I/O.

Backup Results: 2 Clients

Table 29. Backup 2 clients: ADSM V2

File Size	1KB	10KB	100KB	10MB	256MB
MB Transferred	24.4	40.0	100.0	200.0	515.8
Elapsed Time (sec)	640	127	147	260	659
CPU Utilization	49.5	47.4	22.8	15.4	15.5
DASD I/O rate (/sec)	45.5	39.6	31.8	27.2	27.2
DASD Response (msec)	8.1	13.1	13.6	14.9	14.7
DSMSERV CPU Util	86	78	24	9	9
TCPIP CPU Util	10	14	19	19	20
KB/sec	39	323	699	787	802
Total CPU/KB (msec)	25.36	2.93	0.65	0.39	0.39
DASD IOs/KB	1.17	0.12	0.05	0.03	0.03

Note: 9121-480, VM/ESA 2.3.0, 16 Mbit IBM Token Ring, 3172-3, 2 AIX clients on RS/6000 250

Table 30. Backup 2 clients: ADSM V3

File Size	1KB	10KB	100KB	10MB	256MB
MB Transferred	24.5	40.1	100.0	200.0	516.0
Elapsed Time (sec)	224	71	132	254	653
CPU Utilization	49.1	34.3	16.7	13.1	12.9
DASD I/O rate (/sec)	17.1	13.9	7.9	6.3	5.6
DASD Response (msec)	19.7	35.7	56.1	69.9	77.1
DSMSERV CPU Util	66	40	11	4	4
TCPIP CPU Util	28	27	20	20	20
KB/sec	112	578	775	807	809
Total CPU/KB (msec)	8.74	1.19	0.43	0.32	0.32
DASD IOs/KB	0.15	0.02	0.01	0.01	0.01

Note: 9121-480, VM/ESA 2.3.0, 16 Mbit IBM Token Ring, 3172-3, 2 AIX clients on RS/6000 250

Table 31. Backup 2 clients: V3/V2 ratios

File Size	1KB	10KB	100KB	10MB	256MB
MB Transferred	1.00	1.00	1.00	1.00	1.00
Elapsed Time (sec)	0.35	0.56	0.90	0.98	0.99
CPU Utilization	0.99	0.72	0.73	0.85	0.83
DASD I/O rate (/sec)	0.38	0.35	0.25	0.23	0.21
DASD Response (msec)	2.43	2.73	4.13	4.69	5.24
DSMSERV CPU Util	0.77	0.51	0.44	0.48	0.42
TCPIP CPU Util	2.82	1.87	1.08	1.02	1.02
KB/sec	2.88	1.79	1.11	1.03	1.01
Total CPU/KB (msec)	0.34	0.40	0.66	0.83	0.83
DASD IOs/KB	0.13	0.20	0.22	0.23	0.20

Note: 9121-480, VM/ESA 2.3.0, 16 Mbit IBM Token Ring, 3172-3, 2 AIX clients on RS/6000 250

For these measurements, elapsed time was from first client start until last client end. For all measurements, the two clients started and completed their ADSM backups at nearly the same times.

Note that for 1KB file backup, there was a 2.88x throughput increase from V2 for the 2-client measurements as compared to only a 2.26x throughput increase for the 1-client measurements. This greater improvement results from the greatly reduced DSMSERV CPU usage in V3. With V2, DSMSERV CPU utilization with 1 client was 64%. Since DSMSERV CPU utilization cannot exceed 100%, this means that, with V2, a doubling of the backup load by adding a second client cannot be accommodated without a substantial elapsed time increase. With V3, DSMSERV CPU utilization with 1 client was only 42%, resulting in much greater capacity to absorb an increased load without much increase to elapsed time. This effect was much less important for the larger file sizes because the DSMSERV CPU utilizations for these cases were much lower. As a result, the percentage throughput increases relative to V2 for the larger file sizes were quite similar when comparing the 1-client and 2-client measurements.

Generalizing from the above discussion, it is reasonable to conclude that environments characterized by high levels of loading that result in DSMSERV processor and/or I/O contention will tend to experience higher levels of improvement when migrating from V2 to V3 than those shown in this report. This higher degree of improvement results from the large decreases in CPU and I/O usage that were achieved in V3.

Restore Results: 1 Client

Table 32. Restore 1 client: ADSM V2

File Size	1KB	10KB	100KB	10MB	256MB
MB Transferred	12.2	20.0	50.0	100.0	258.0
Elapsed Time (sec)	555	111	204	270	695
CPU Utilization	31.9	30.4	13.1	14.2	12.6
DASD I/O rate (/sec)	4.6	6.5	6.9	9.4	9.4
DASD Response (msec)	18.3	20.5	21.5	22.5	22.9
DSMSERV CPU Util	60	51	12	9	5
TCPIP CPU Util	2	8	12	16	17
KB/sec	22.5	184.8	251.4	378.7	379.9
Total CPU/KB (msec)	28.32	3.29	1.04	0.75	0.66
DASD IOs/KB	0.20	0.04	0.03	0.02	0.02

Note: 9121-480, VM/ESA 2.3.0, 16 Mbit IBM Token Ring, 3172-3, AIX 4.2.1 client on RS/6000 250

Table 33. Restore 1 client: ADSM V3

File Size	1KB	10KB	100KB	10MB	256MB
MB Transferred	12.27	20.05	50.02	100	257.94
Elapsed Time (sec)	504.3	96.7	206.2	279.9	707.4
CPU Utilization	17.0	24.5	10.2	8.9	8.4
DASD I/O rate (/sec)	2.7	6.7	6.6	9.1	9.2
DASD Response (msec)	14	19.5	21	22.1	22.4
DSMSERV CPU Util	30	40	9	3	1
TCPIP CPU Util	1	6	9	12	12
KB/sec	25	212	248	366	373
Total CPU/KB (msec)	13.65	2.31	0.82	0.49	0.45
DASD IOs/KB	0.11	0.03	0.03	0.02	0.02

Note: 9121-480, VM/ESA 2.3.0, 16 Mbit IBM Token Ring, 3172-3, AIX 4.2.1 client on RS/6000 250

Table 34. Restore 1 client: V3/V2 ratios

File Size	1KB	10KB	100KB	10MB	256MB
MB Transferred	1.01	1.00	1.00	1.00	1.00
Elapsed Time (sec)	0.91	0.87	1.01	1.04	1.02
CPU Utilization	0.53	0.81	0.78	0.63	0.67
DASD I/O rate (/sec)	0.59	1.03	0.96	0.97	0.98
DASD Response (msec)	0.77	0.95	0.98	0.98	0.98
DSMSERV CPU Util	0.51	0.79	0.79	0.28	0.27
TCPIP CPU Util	0.84	0.80	0.72	0.76	0.73
KB/sec	1.11	1.15	0.99	0.97	0.98
Total CPU/KB (msec)	0.48	0.70	0.79	0.65	0.68
DASD IOs/KB	0.53	0.90	0.97	1.00	1.00

Note: 9121-480, VM/ESA 2.3.0, 16 Mbit IBM Token Ring, 3172-3, AIX 4.2.1 client on RS/6000 250

Appendix A. Workloads

The workloads that were used to evaluate VM/ESA 2.4.0 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 35 shows the search-order characteristics of the two environments used for measurements discussed in this document.

Table 35. 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 71 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 36 on page 60.

CMS-Intensive (FS8F)

Table 36. FS8F workload script summary

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

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

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

LOGESA: Initialization Script

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

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

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

ASM617F: Assemble (HLASM) and Run

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

ASM627F: Assemble (F-Assembler) and Run

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

XED117F: Edit a VS BASIC Program

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

XED127F: Edit a VS BASIC Program

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

XED137F: Edit a COBOL Program

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

XED147F: Edit a COBOL Program

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

COB217F: Compile a COBOL Program

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

COB417F: Run a COBOL Program

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

FOR217F: Compile 6 VS FORTRAN Programs

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

FOR417F: Run 2 FORTRAN Programs

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

PRD517F: Productivity Aids Session

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

DCF517F: Edit and SCRIPT a File

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

PLI317F: Edit and Compile a PL/I Optimizer Program

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

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

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

WND517F: Use Windows

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

WND517FL: Use Windows with LOGON, LOGOFF

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

HLP517F: Use HELP and Miscellaneous Commands

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

VSE Guest (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:

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

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:

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

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:

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

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:

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

First_Response is the average first response given in the RSPRPT section of the TPNS reports. *Total_Transmits* is the total TPNS transmits and *TPNS_Time* is the run interval log time found in the Summary of Elapsed Time and Times Executed section of the TPNS reports.

AVG LAST (T). The average response time in seconds for the last response to the screen. If there is more than one TPNS this is calculated by:

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

Last_Response is the average last response given in the RSPRPT section of the TPNS reports. *Total_Transmits* is the total TPNS transmits and *TPNS_Time* is the run interval log time found in the Summary of Elapsed Time and Times Executed section of the TPNS reports.

Glossary of Performance Terms

AVG Lock Wait Time (ms). The average time it takes for an SFS lock conflict to be resolved. This is calculated by:

$$\frac{\sum_{f \in \text{filepools}} \frac{Lock\_Wait\_Time_f}{SFSTIME_f}}{\sum_{f \in \text{filepools}} \frac{Total\_Lock\_Conflicts_f}{SFSTIME_f}}$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Agent\_Holding\_Time_f}{SFSTIME_f}}{\sum_{f \in \text{filepools}} \frac{Begin\_LUWs_f}{SFSTIME_f}}$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Total\_BIO\_Request\_Time_f}{SFSTIME_f}}{\sum_{f \in \text{filepools}} \frac{Total\_BIO\_Requests_f}{SFSTIME_f}}$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Total\_DASD\_Block\_Transfers_f}{SFSTIME_f}}{\sum_{f \in \text{filepools}} \frac{Total\_BIO\_Requests_f}{SFSTIME_f}}$$

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:

$$\frac{\sum_{f \in \text{filepools}} \frac{Total\_DASD\_Block\_Transfers_f}{SFSTIME_f}}{\sum_{f \in \text{filepools}} \frac{Total\_IO\_Requests_f}{SFSTIME_f}}$$

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:

$$\frac{1}{10} \times \sum_{f \in \text{filepools}} \frac{Checkpoint\_Time_f}{SFSTIME_f}$$

Checkpoint_Time is from the QUERY FILEPOOL STATUS command.

Checkpoints Taken (delta). The number of checkpoints taken by all file pools on the system. This is calculated by:

$$\sum_{f \in \text{filepools}} Checkpoints\_Taken_f$$

Checkpoints_Taken is from the QUERY FILEPOOL STATUS command.

CMS BLOCKSIZE. The block size, in bytes, of the users′ CMS minidisks.

Command. In the context of reporting performance results, any user interaction with the system being measured.

CP/CMD. For the FS7F, FS8F, and VSECICS workloads, this is the average amount of CP processor time used per command in milliseconds. For the PACE workload, this is the average CP processor time per job in seconds. This is calculated by:

For the FS7F, FS8F, and VSECICS workloads:

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

For the PACE workload:

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

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

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

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

For the PACE workload:

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

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

For all workloads running on 9121 and 9021 processors:

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

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

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

TCPU is Total CPU busy seconds, *VCPU* is Virtual CPU seconds, and *V_Time* is the VMPRF time interval obtained from the Resource Utilization by User Class section of the VMPRF report.

CPU PCT BUSY (V). CPU Percent Busy. The percentage of total available processor time used by the designated virtual machine. Total available processor time is the sum of online time for all processors and represents total processor capacity (not processor usage).

This is the from the CPU Pct field in the VMPRF USER_RESOURCE_USER report.

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

CPU UTIL (V). The percentage of 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 X′04′ instructions used per command. DIAGNOSE code X′04′ is the privilege class C and E CP function call to examine real storage. This is a product-sensitive programming interface. This is calculated by:

$$\frac{DIAG\_04}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_04 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 08/CMD. The number of DIAGNOSE code X′08′ instructions used per command. DIAGNOSE code X′08′ is the CP function call to issue CP commands from an application. This is calculated by:

$$\frac{DIAG\_08}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_08 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 0C/CMD. The number of DIAGNOSE code X′0C′ instructions used per command. DIAGNOSE code X′0C′ is the CP function call to obtain the time of day, virtual CPU time used by the virtual machine, and total CPU time used by the virtual machine. This is calculated by:

$$\frac{DIAG\_0C}{\text{RTM\_Time} \times \text{ETR (T)}}$$

DIAG_0C is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 10/CMD. The number of DIAGNOSE code X′10′ instructions used per command. DIAGNOSE code X′10′ is the CP function call to release pages of virtual storage. This is calculated by:

$$\frac{DIAG\_10}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_10 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 14/CMD. The number of DIAGNOSE code X′14′ instructions used per command. DIAGNOSE code X′14′ is the CP function call to perform virtual spool I/O. This is calculated by:

$$\frac{DIAG\_14}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_14 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 58/CMD. The number of DIAGNOSE code X′58′ instructions used per command. DIAGNOSE code X′58′ is the CP function call that enables a virtual machine to communicate with 3270 virtual consoles. This is calculated by:

$$\frac{DIAG\_58}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_58 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 7C/CMD. The number of DIAGNOSE code X′7C′ instructions used per command. DIAGNOSE code X′7C′, 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 X′98′ instructions used per command. This allows a specified virtual machine to lock and unlock virtual pages and to run its own channel program. This is calculated by:

$$\frac{DIAG\_98}{\text{RTM\_Time} \times \text{ETR (T)}}$$

*DIAG*_98 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 98/CMD (V) VTAM Servers. See DIAG 98/CMD for a description of this instruction. This represents the sum of all DIAGNOSE code X′98′ instructions per command for all VTAM and VSCS servers. This is calculated by:

$$\frac{DIAG\_98\_VTAM + DIAG\_98\_VSCS}{\text{ETR (T)}}$$

*DIAG*_98_*VTAM* and *DIAG*_98_*VSCS* are taken from the VMPRF Virtual Machine Communication by User Class report for the VTAM and VSCS server classes respectively.

DIAG A4/CMD. The number of DIAGNOSE code X′A4′ instructions used per command. DIAGNOSE code X′A4′ is the CP function call that supports synchronous I/O to supported DASD. This is calculated by:

$$\frac{DIAG\_A4}{\text{RTM\_Time} \times \text{ETR (T)}}$$

DIAG_A4 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG A8/CMD. The number of DIAGNOSE code X′A8′ instructions used per command. DIAGNOSE code X′A8′ is the CP function call that supports synchronous general I/O to fully supported devices. This is calculated by:

$$\frac{DIAG\_A8}{RTM\_Time \times ETR\ (T)}$$

DIAG_A8 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

DIAG 214/CMD. The number of DIAGNOSE code X′214′ instructions used per command. DIAGNOSE code X′214′ is used by the Pending Page Release function. This is calculated by:

$$\frac{DIAG\_214}{RTM\_Time \times ETR\ (T)}$$

*DIAG*_214 is taken from the TOTALCNT column on the RTM PRIVOPS screen. RTM_Time is the total RTM time interval.

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

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

*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 X′270′ instructions used per command. DIAGNOSE code X′270′ 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 X′0C′ 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:

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

*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:

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

For the PACE workload:

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

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).

$$\frac{1}{Num\_Entries} \times \sum_{t \in SCLOG\ entries} Q0_t + Q0L_t + Q1_t + Q1L_t + Q2_t + Q2L_t + Q3_t + Q3L_t$$

$Q0_t$, $Q0L_t$ .. 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).

$$\frac{1}{Num\_Entries} \times \sum_{t \in SCLOG\ entries} E0_t + E0L_t + E1_t + E1L_t + E2_t + E2L_t + E3_t + E3L_t$$

$E0_t$, $E0L_t$ .. 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:

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

For the PACE workload, this is calculated by:

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

EMUL/CMD (H). See EMUL/CMD. This is the hardware based measurement.

For the 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:

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

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:

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

Total_BIO_Request_Time is from the QUERY FILEPOOL STATUS command.

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

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

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:

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

ITR (H). See ITR. This is the hardware based measure. In this case, ITR is measured in external commands per unit of processor busy time. For the 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:

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

For the PACE workloads:

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

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:

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

ITRR. Internal Throughput Rate Ratio. This is the RTM based ITR normalized to a specific run. This is calculated by:

$$\frac{ITR}{ITR_1}$$

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:

$$\frac{ITR\ (H)}{ITR\ (H)_1}$$

$ITR\ (H)_1$ 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:

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

$ITR\ (V)_1$ 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:

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

For the PACE workload:

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

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

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

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

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

Path length. See Instruction Path Length

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

For the FS7F, FS8F, and VSECICS workloads:

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

For the PACE workload:

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

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

For the FS7F, FS8F, and VSECICS workloads:

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

For the PACE workload:

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

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

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

PGBLPGS. The number of system pageable pages available.

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

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

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

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

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

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

For the FS7F, FS8F, and VSECICS workloads:

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

For the PACE workload:

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

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

PRIVOPS (Privileged Operations). See Privileged Operation.

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

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

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

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

PSU. Product Service Upgrade

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

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

RAID. Redundant array of independent DASD.

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

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

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

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

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

RESERVE. See SET RESERVED

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

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:

$$\sum_{f \in \text{filepools}} Rollback\_Requests_f$$

Rollback_Requests is from the QUERY FILEPOOL STATUS command.

Rollbacks Due to Deadlock (delta). The total number of LUW rollbacks due to deadlock that occurred during the measurement interval over all production file pools. A rollback occurs whenever a deadlock condition cannot be resolved by the SFS server. This is calculated by:

$$\sum_{f \in \text{filepools}} Rollbacks\_Due\_to\_Deadlock_f$$

Rollbacks_Due_to_Deadlock is from the QUERY FILEPOOL STATUS command.

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

RSU. Recommend Service Upgrade

RTM. Realtime Monitor. A licensed program realtime monitor and diagnostic tool for performance monitoring, analysis, and problem solving.

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

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

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

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

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

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

Checkpoints_Taken is from the QUERY FILEPOOL STATUS command.

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

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

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

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

Filepool_Request_Service_Time is from the QUERY FILEPOOL STATUS command.

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

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

SHRPGS. The number of shared frames currently resident.

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

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

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

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

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

For the FS7F, FS8F, and VSECICS workloads:

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

For the PACE workload:

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

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

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

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

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

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:

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