VM/ESA RELEASE 1.0 PERFORMANCE REPORT

January 28, 1991

VM Performance
Departments G21/G25
Endicott, NY

VM/XA Performance
Departments 46DA/46EA
Kingston, NY

VTAM Performance
Department A37A
Raleigh, NC

IBM Internal Use Only

Preface

Performance data contained within this report represent specific environments created and measured in Endicott by the VM Performance Department and in Kingston by the VM/XA Performance Department and the Raleigh VTAM Performance Department. These environments contained I/O configurations, programming products, run procedures, options, microcode assists, and workloads designed for specific processors. The resulting data may or may not be comparable to other existing performance data due to subtle differences in any of the aforementioned.

Abstract

VM/ESA 1.0 Performance Report provides information for VM/ESA 1.0 running on 9370, 4381, and 3090 processors.

The intent of this report is to provide performance and tuning information based on the results of the VM/ESA 1.0 performance test conducted jointly by the Endicott, Kingston, and Raleigh programming laboratories.

Discussion centers on two basic environments, high end VM/ESA performance and performance of the VM/ESA 370 Feature. The high end VM/ESA performance section deals with regression for the CMS Intensive environment, Migration, Processor Capacity for the CMS Intensive workload, VM Connectivity, single VTAM, and running MVS Guests. The VM/ESA 370 Feature section covers regression, of the CMS Intensive and SFS environments, the CMS Shared File System, VM Connectivity, Coordinated Resource Recovery, GCS, and running with more than 16M of real storage.

All measurements were made by the personnel listed in the acknowledgements section.

Acknowledgements

The following personnel contributed to this report:

Endicott Programming Laboratory

- **M. S. Bidwell**
- **W. J. Bitner**
- **C. T. Bradley**
- **W. G. Ernsberger**
- **G. S. Gasper**
- **G. A. Hine**
- **R. R. Kaskan**
- **A. B. Lebovitz**
- **S. P. Lyons**
- **S. T. Marcotte**
- **V. E. Meredith**
- **E. J. Strom**
- **M. S. Thomas**
- **J. L. Thrall**
- **J. S. Tingley**
- **A. M. Ward**
- **R. C. Zinna**

Kingston Programming Laboratory

- **R. D. Deese**
- **C. L. Dewar**
- **E. M. Digan**
- **J. F. Eschmann**
- **T. J. Fenton**
- **M. Ferrell**
- **B. A. Hall**
- **B. C. Holmquist**
- **R. A. Klein**
- **S. S. Lee**
- **F. D. Lewis**
- **J. P. Spano**
- **K. P. Tsai**
- **D. A. Vastola**

Raleigh Programming Laboratory

- **C. J. Ames**
- **J. T. Jennings**
- **R. A. Holton**
- **K. H. Menditto**

Table of Contents

List of Illustrations

List of Tables

1.0 Introduction

This report describes performance characteristics of VM/ESA 1.0. The measurements within the report fall into two categories - system regression and new function. System regression measurements display the effects of design changes, service, and new code which enter VM over time. Since VM/ESA 1.0 (ESA 1.0) incorporates VM/SP 6 (SP 6), VM/HPO 5.0 (HPO 5.0), and VM/XA SP 2.1 (XA 2.1), regression measurements in this report compare VM/ESA 1.0 to VM/SP 6, VM/HPO 5.0, or VM/XA SP 2.1, as applicable. There are also comparisons to VM/XA SP 2.0 for informational purposes. The new function measurements describe the performance of new items appearing in VM/ESA 1.0.

An additional distinction can be applied to the measurement data in this report. The data separates into two sections - high end measurements, which focus primarily on large processor environments, and 370 Feature measurements which evaluate the VM/ESA 1.0 370 Feature. Processor sizes range from 9370 and 4381 processors to a 3090-600, giving a broad range of CPU speeds and processor sizes. The 9370 and 4381 systems were built from a System Offering base. VM/SP Release 6 on these processors refers to VM/SP Release 6 at Service Level 601. The 3090 system was built from a VM/XA SP Release 2.1 base. VM/XA SP Release 2.1 refers to VM/XA SP Release 2.1 at Service Level 0.

All measurements in this report were made on a VM/ESA 1.0 ISD-equivalent driver unless otherwise noted.

2.0 General Observations

2.1 VM/ESA 1.0

2.1.1 Comparisons to Other Workloads

All of the CMS regression measurement data published in this report was obtained with the FS7B workload in its various versions which allows for comparisons of the EDF and SFS file systems, and for one workload throughout the processor range. An in-depth description of FS7B is located in "CMS Interactive (FS7B)" on page 163. FS7B was chosen because it was the most customer representative CMS intensive workload, it exercised the most functions, it easily allows for switching between EDF and SFS, and it gives the Development organization incentive to improve many areas of the system. Given that no one workload can represent all of CMS's varied customers, FS7B provides the best overall measure of the strengths and weaknesses of the release. VM/ESA 1.0 compares very favorably to its predecessors when measured with the FS7B workload. It is expected that most CMS intensive installations will realize equivalence with their previous release of VM up through the range of improvement measured and reported here. Significant improvements have been made in CMS to COPYFILE, Assemble (CMSSTOR), EXECIO within REXX, FST Lookup and page aligned buffer management to better exploit minidisk caching. Users of these functions will see the largest benefit.

The LSPR (Large System Performance Report) workloads (PD3 and HT4), which were developed for the purpose of comparing processors, were also measured against VM/XA SP 2.1 and VM/ESA 1.0. Since the LSPR workloads do not exercise a broad scope of functions nor do they invoke new functions such as windows and SFS, that data is not emphasized in this report. That data will be available in a separate report, but is worth a brief mention here since it tells a slightly different story than the FS7B results. VM/ESA 1.0 realized a 2-4% ITR degradation compared to VM/XA SP 2.1 when measured with the LSPR workloads. Investigating the reasons behind these results identified the following areas of concern. Each of the PD3 scripts and five of the twenty HT4 scripts executes an IPL CMS and because of the inevitable growth of CMS this command has degraded. However, it is felt that the average user does not IPL CMS this frequently and that this is an anomaly of the workload done to ensure repeatability. Another area which exhibited growth is the issue of a PF11 within FILELIST, but only for one specific case. In VM/ESA 1.0, an additional REXX exec is called by PF11 to execute XEDIT when FILELIST is entered with a filemode of "*". This is due to SFS implications because both directories and EDF files may be present when FILELIST is issued with filemode "*". Again, this particular variation of the FILELIST command is heavily issued in the LSPR workloads. Finally, there is growth in the execution of a COBOL program. All of these degraders are legitimate and in the interest of quality, fixes will be investigated for the degraded areas exposed by LSPR. However, since LSPR is a different workload than FS7B, it takes less advantage of the improvements mentioned above, therefore an overall slight degradation was observed.

The majority of CMS intensive installations will realize results closer to what was measured with FS7B. Since the LSPR workloads are ingrained in the performance community, discussion of them was warranted here.

2.1.2 Regression

The purpose of this section is to summarize performance differences that have occurred due to the addition of new function and service since the introduction of VM/HPO 5.0 and VM/XA SP 2.1. The workloads executed were not changed explicitly to use any of the new ESA 1.0 functions. These comparisons can be used as an indicator of what might be expected with similar workloads when up-leveling to this new level of VM.

Some of the changes that were made in VM/ESA 1.0 that have the potential to affect the performance are:

In CP, Pageable Page Tables support was added. This benefits users with large virtual machines. It allows real storage, that contain CP control blocks that map virtual memory that is not in use, to be made available for other uses. Virtual machine size relief, allowing virtual machine sizes greater than 999MB, and up to 2047MB, was also added.

In CMS, the nucleus area has grown. The storage referenced for moderate to large accessed disks has been reduced. The Processor Busy Time (PBT) used by CMS has also changed. Several additional changes were made, with the most notable ones being reduced processor busy time for EXECIO when called from REXX, COPYFILE improvements on a per record basis, Assembler time reductions due to a reduction in CMSSTOR macro expansion, and an increase in the IPL CMS time. The EXECIO improvement is also responsible for the large reduction in the PRIVOP rate. The SPKA instructions were eliminated from the linkage path between the REXX and EXECIO modules. The virtual file I/O rate showed a slight reduction mainly due to a reduction in I/Os for the Assembler when expanding the CMSSTOR macro. The performance effect of these changes is workload dependent. The improvements will show best in those environments that have many assemblies.

GCS/XA has been changed to allow its virtual machines to use 31 bit addressing. GCS, in supervisor state, still runs in a 24 bit address space.

VTAM has been changed to allow its virtual machines to use 31 bit addressing.

The CMS Intensive environments that were used for the regression measurements were:

Comparing VM/ESA 1.0 with VM/HPO 5.0 on a 3090-200J.

Comparing VM/ESA 1.0 with VM/XA SP 2.1 on a 3090-200J.

Comparing VM/ESA 1.0 with VM/XA SP 2.0 on a 3090-600J.

Comparing VM/ESA 1.0 with VM/XA SP 2.1 on a 3090-600J.

Comparing VM/ESA 1.0 with VM/HPO 5.0 on a 3090-200J

The measurements on the 3090-200J comparing VM/ESA 1.0 to VM/HPO 5.0 were done in the following manner. First, a measurement was done to determine how many users, running the CMS intensive workload (FS7B0R) in 370 mode, it took to run VM/HPO 5.0 at about a 90% processor utilization on the 3090-200J. To obtain a 90% processor utilization, the think time had to be reduced to 15 seconds. With the 30 second think time, used for the rest of the CMS intensive measurements, HPO 5.0 became storage constrained and a valid 90% processor utilization measurement could not be made. This same number of users (1140) and think time (15 seconds) were then used for a VM/ESA 1.0 measurement. All CMS users were running in 370 Mode.

The results of these measurements show that the VM/ESA 1.0 response times and ETR were better than and ITR was about the same as the VM/HPO 5.0 release. The following figures summarize these results and section "CMS Intensive Measurements" on page 38 contains the details of the measurements.

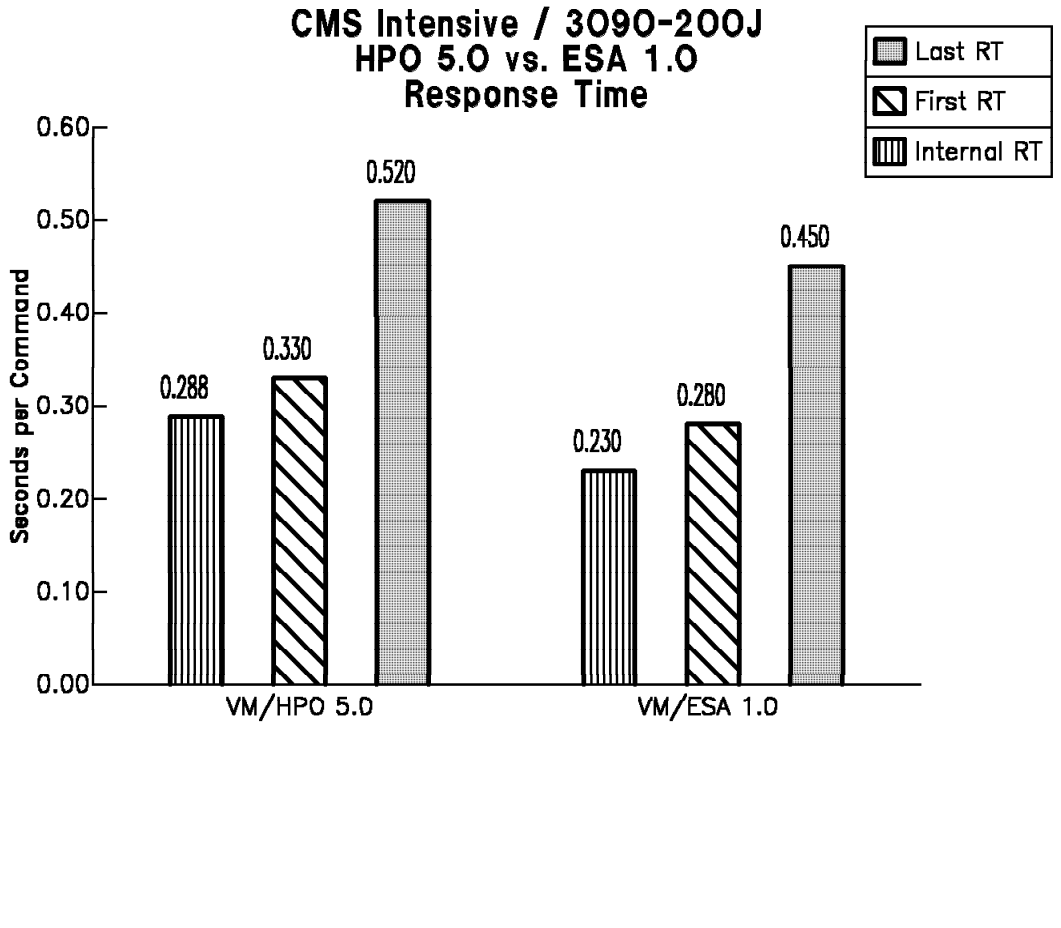

Figure 1. ESA 1.0 vs. HPO 5.0 -- Response Time Comparisons: 1140 CMS Intensive Users with 370 Mode CMS Virtual Machines

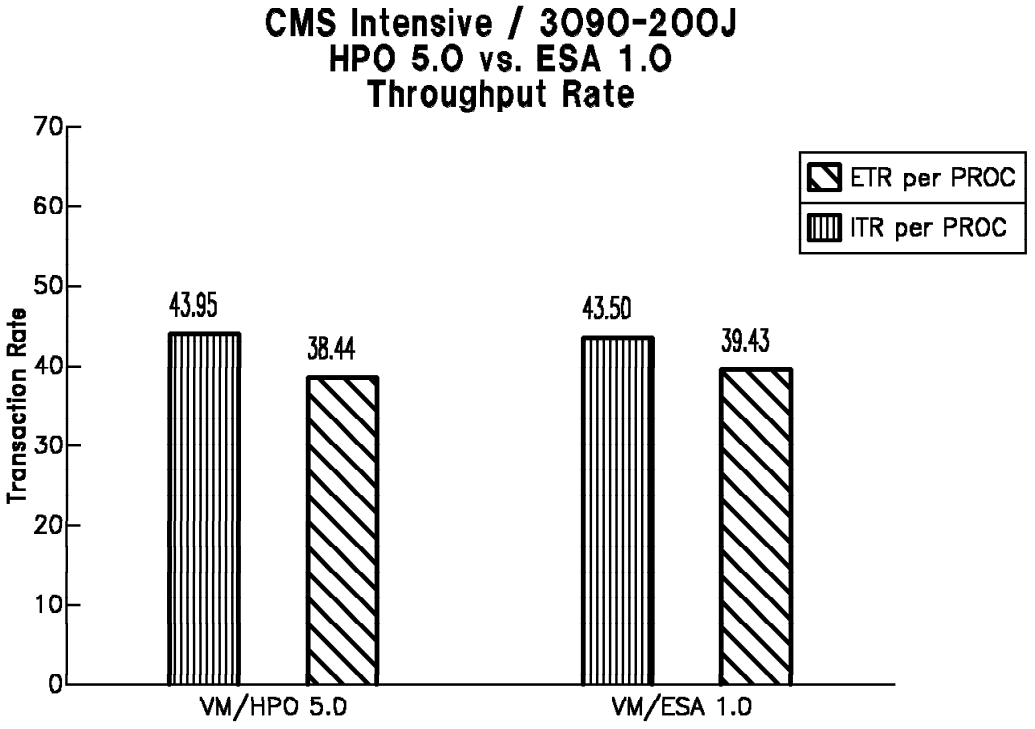

Figure 2. ESA 1.0 vs. HPO 5.0 -- Throughput: 1140 CMS Intensive Users with 370 Mode CMS Virtual Machines

Comparing VM/ESA 1.0 with VM/XA SP 2.1 on a 3090-200J

The measurements on the 3090-200J comparing VM/ESA 1.0 to VM/XA SP 2.1 were done in the following manner. First, a measurement was done to determine how many users, running the CMS intensive workload (FS7B0R) in 370 mode, it took to run VM/XA SP 2.1 at a 90% processor utilization on the 3090-200J. This same number of users (2000) was then used for a VM/ESA 1.0 measurement. All CMS users were running in 370 Mode.

The results of these measurements show that VM/ESA 1.0 has better overall response times and throughput than the VM/XA SP 2.1 release. The following figures summarize these results and section "3090-200J / 256M / VM/XA SP 2.1" on page 43 contains the details of the measurements.

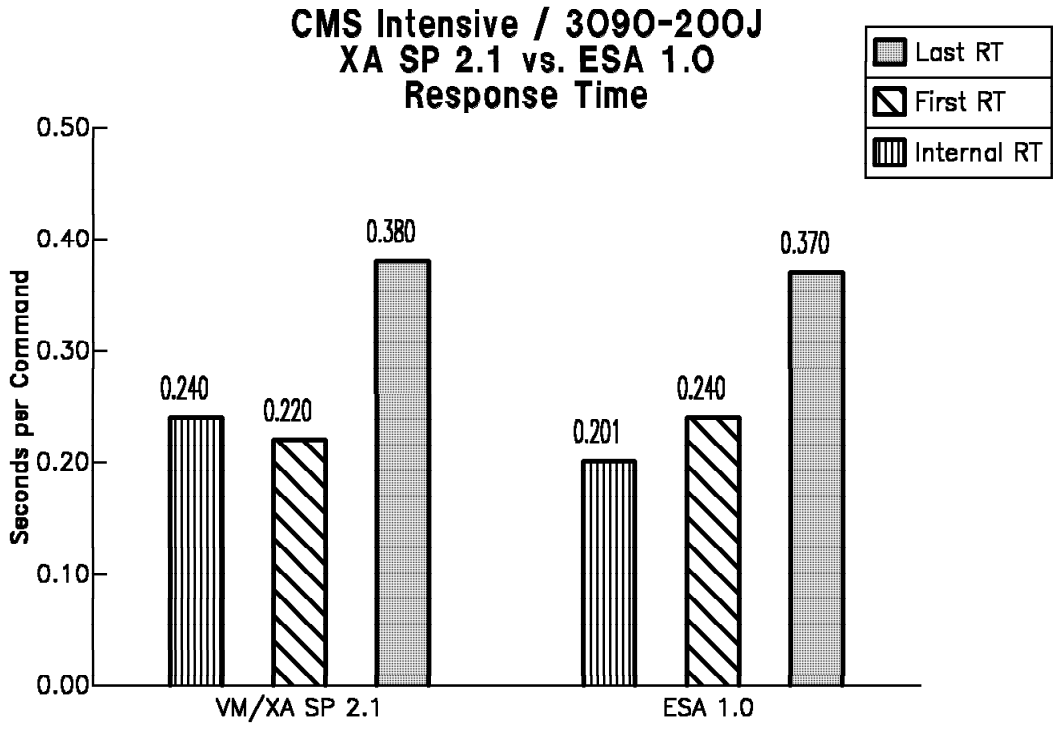

Figure 3. Response Time Comparison with 2000 CMS Intensive Users: 370 Mode CMS Virtual Machines

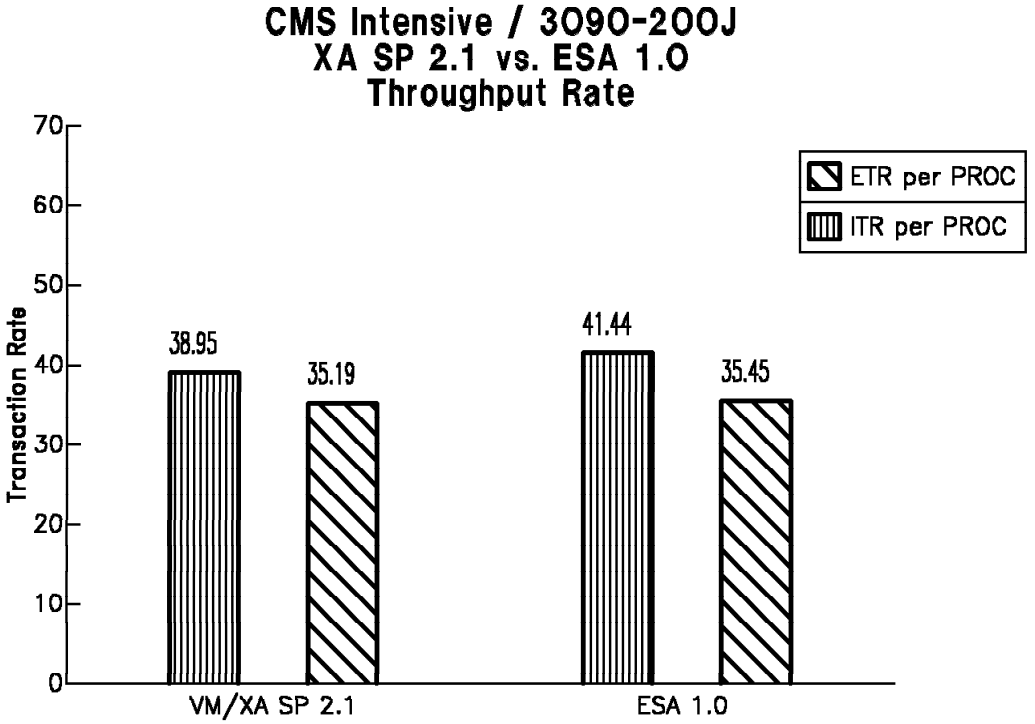

Figure 4. Throughput Comparison with 2000 CMS Intensive Users: 370 Mode CMS Virtual Machines

Comparing VM/ESA 1.0 with VM/XA SP 2.0 and VM/XA SP 2.1 on a 3090-600J

The measurements on the 3090-600J comparing VM/ESA 1.0 to VM/XA SP 2.0 and VM/XA SP 2.1 were done in the following manner. First, a measurement was done to determine how many users, running the CMS intensive workload (FS7B0R) in 370 mode, it took to run VM/XA SP 2.0 at a 90% processor utilization on the 3090-600J. This same number of users (5400) was then used for VM/XA SP 2.1 and VM/ESA 1.0 measurements. All CMS users were running in 370 Mode.

The results of these measurements show that VM/ESA 1.0 has better response times and throughput than either of the XA/SP releases. The following figures summarize these results and sections "3090-600J / 512M / VM/XA SP 2.0" on page 41 and "3090-600J / 512M / VM/XA SP 2.1" on page 45 contain the details of the measurements.

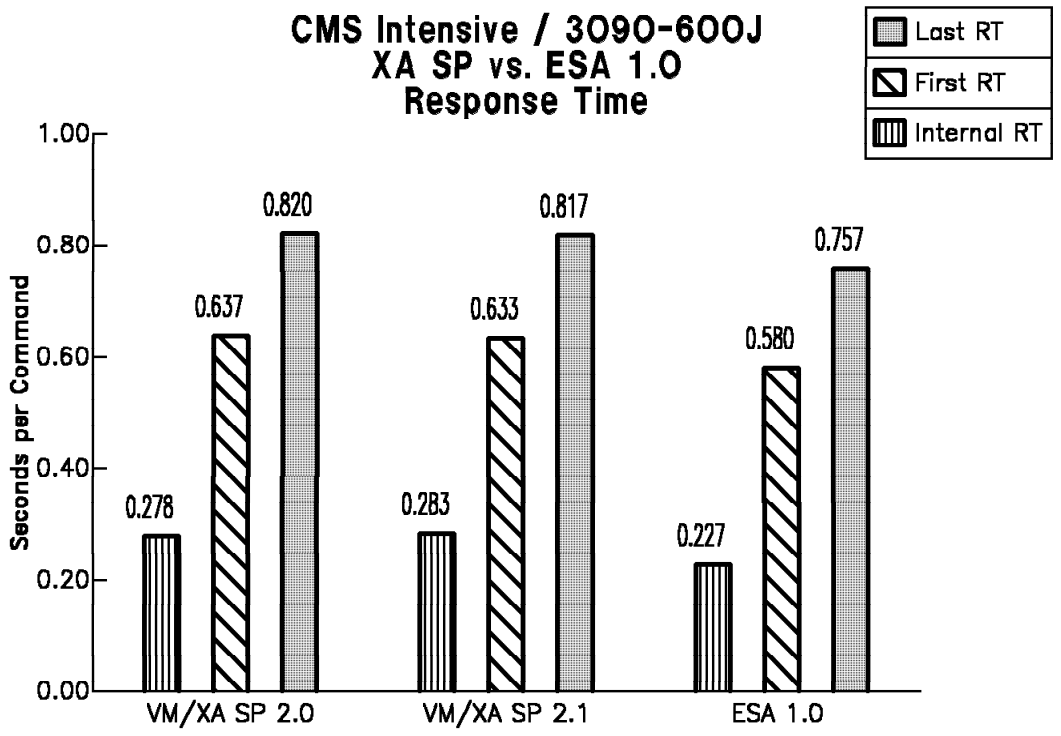

Figure 5. Response Time Comparison with 5400 CMS Intensive Users: 370 Mode CMS Virtual Machines

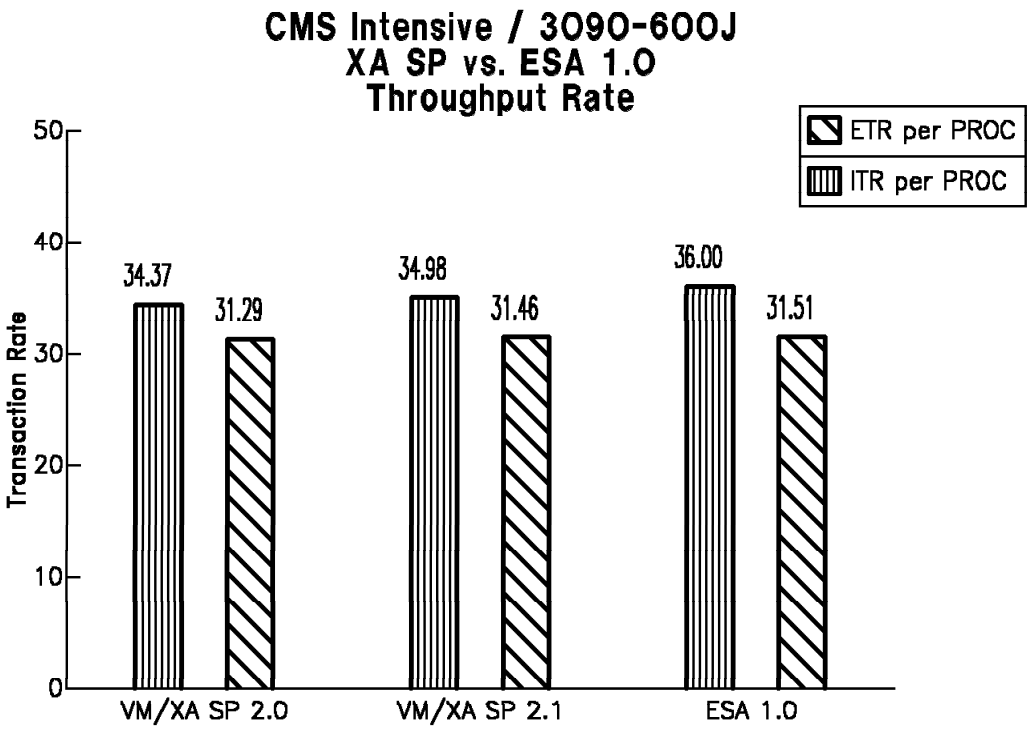

Figure 6. Throughput Comparison with 5400 CMS Intensive Users: 370 Mode CMS Virtual Machines

2.1.3 Comparison of 370 and XA Modes

A set of CMS intensive measurements was taken to compare the performance of users running 370 mode virtual machines with those running XA mode virtual machines. The measurements were done on the 3090-600J with 5400 users, giving a target processor utilization of about 90%.

The measurements were done in the following manner. First, a measurement was done to determine how many users, running the CMS intensive workload (FS7B0R) in 370 mode, it took to run the 3090-600J at a 90% processor utilization. This same workload was then run with the same number of users (5400) in XA mode.

A summary of the results is that for XA Mode:

internal response time was about 0.03 seconds faster
external response time took less than 0.25 seconds longer
ITR was reduced by less than 3%
ETR was reduced by less than 1%
processor usage increased about 2%.

The following figures summarize these results and section “370 to XA Measurements” on page 47 contains the details of the measurements.

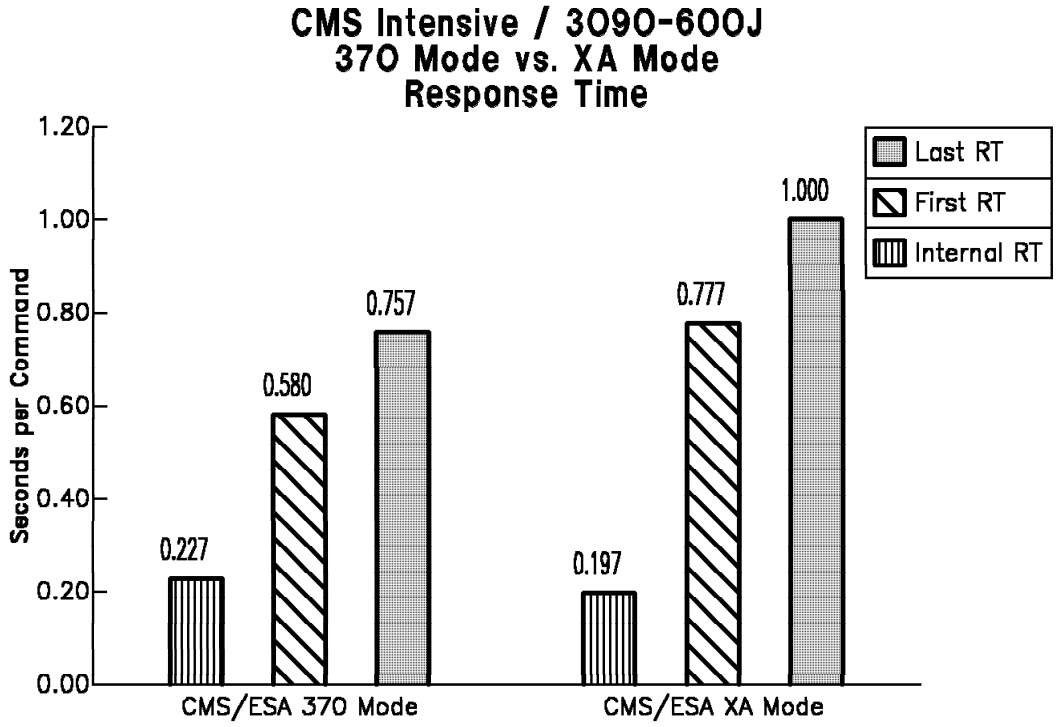

Figure 7. Response Time Comparison of 370 Mode vs. XA Mode CMS Users

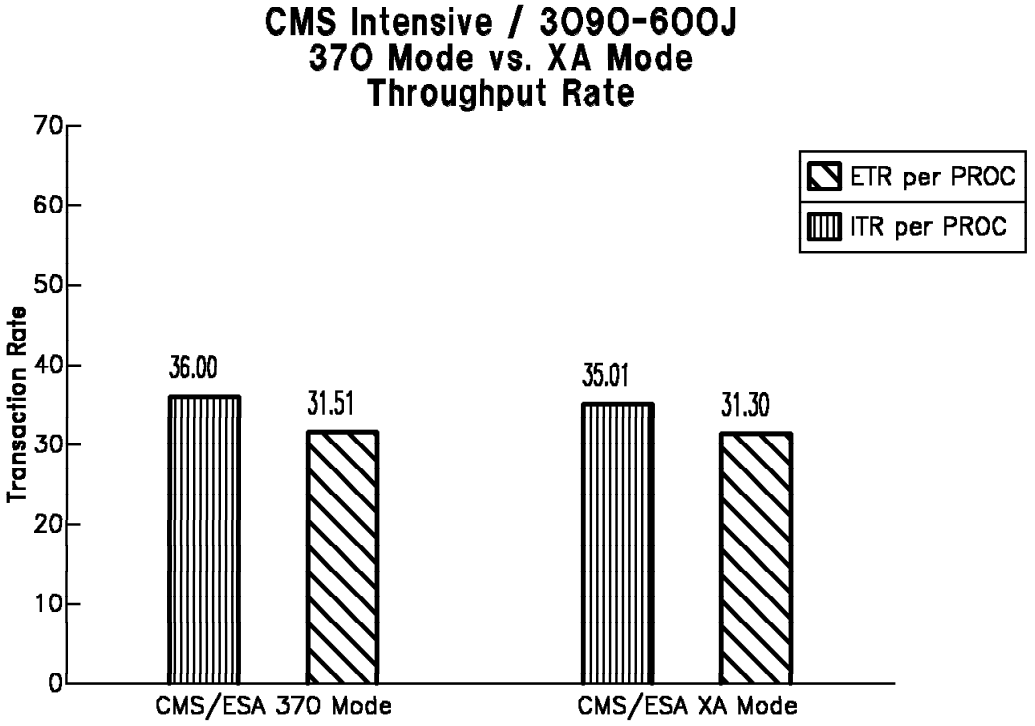

Figure 8. Throughput Comparison of 370 Mode vs. XA Mode CMS Users

2.1.4 EDF and SFS Comparisons

Measurements were obtained on ESA 1.0 to compare the 0% SFS (EDF) workload with the 35% SFS workload (all user files are in SFS). For a more complete description of these workloads, see "Appendix C. Workloads" on page 163.

The measurements show that SFS requires more processor time and real storage than EDF when performing similar tasks. However, SFS has significant advantages over minidisk, such as:

superior DASD space management and efficiency
hierarchical directories
full file-level sharing
file-level Security
file aliasing
distributed (remote) file access
high-level language callable API
superior data integrity
improved file referencing

These large system measurements and the ESA 1.0 370 Feature measurements both show very similar SFS performance impacts, indicating that SFS performance scales to large systems. Large systems, however, will generally require multiple file pools. For more information about the large system measurements see "EDF and SFS Comparisons" on page 50. For more information about the ESA 1.0 370 Feature measurements see "EDF and SFS Comparisons" on page 101.

Response times for EDF and SFS are similar with similar processor utilization.

The minidisk cache (MDC) is equally effective for EDF and SFS at reducing DASD read I/Os. Maintenance of MDC requires less processing in the SFS case.

Two views are provided:

1. Equal Utilization: EDF and SFS measurements are compared at the same processor utilization.

   - A 16% reduction in the number of users was required to hold CPU utilization constant.
   - Processor busy time per command increased 14%.
   - Real storage contention (PAGE/CMD and XSTOR/CMD) increased 11%.
   - Virtual file I/Os per command increased 7%.
   - Total internal response time per command was 0.44 seconds for EDF and 0.41 seconds for SFS.

2. Equal Users: EDF and SFS measurements are compared at the same number of users.

   - Processor busy time per command increased 17%.
   - Real storage contention (PAGE/CMD and XSTOR/CMD) increased 64%.
   - Virtual file I/Os per command increased 9%.
   - Total internal response time per command increased from 0.22 seconds for EDF to 0.41 seconds for SFS.

Note: A Coordinated Resource Recovery (CRR) server did exist for these measurements, but in this environment the recovery server is not involved in mainline processing. If the recovery server had not been running, limp mode could increase processor requirements by as high as 40%. To have acceptable performance you must have a CRR server running, even in regression environments where you are not using CRR.

2.1.5 Connectivity

In this document *Connectivity* refers to APPC communications between virtual machines. The following three types of APPC connections between virtual machines are considered:

APPC/VM communications between virtual machines on the same system, using the APPCVM macro.

TSAF communications between virtual machines on two distinct systems in a TSAF collection, using the TSAF (Transparent Services Access Facility) component of VM.

AVS communications between virtual machines on two distinct collections, using VTAM and the AVS (APPC/VM VTAM Support) component of VM.

The same tool is used to drive all these connections - the Connectivity Driver, described in "Connectivity Driver Tool" on page 176. This tool consists of *user* and *resource* programs running in separate virtual machines. In general, users make requests of resources and resources send data to users. Resources can handle many users at once, but each user can connect to only one resource at a time.

Two of these three connectivity environments were measured on ESA 1.0:

- APPC/VM performance compared to HPO 5.0
- AVS performance.

2.1.5.1 APPC/VM

The performance of APPC/VM on ESA 1.0 was compared to an equivalent workload on HPO 5.0. HPO 5.0 was chosen as the most likely migration path for high-end users of APPC. Both single-thread (one user, one resource) and multi-user (many users, two resources) environments were measured. Measurements were made on a 3090-300J with one, or in some cases two, processors offline.

In terms of processor time per transaction, APPC/VM is slower on ESA 1.0 than on HPO 5.0, when the environments are made as equal as possible. Figure 9 on page 15 and Figure 10 on page 16 show the degradation from HPO 5.0 when running the single-thread tests on 1 and 2 processors respectively. The benchmarks consisted of running many consecutive APPC transfers of 512, 4096 and 102400 bytes from a resource machine to a user machine. These measurements are described in detail in "APPC/VM Single-user" on page 59.

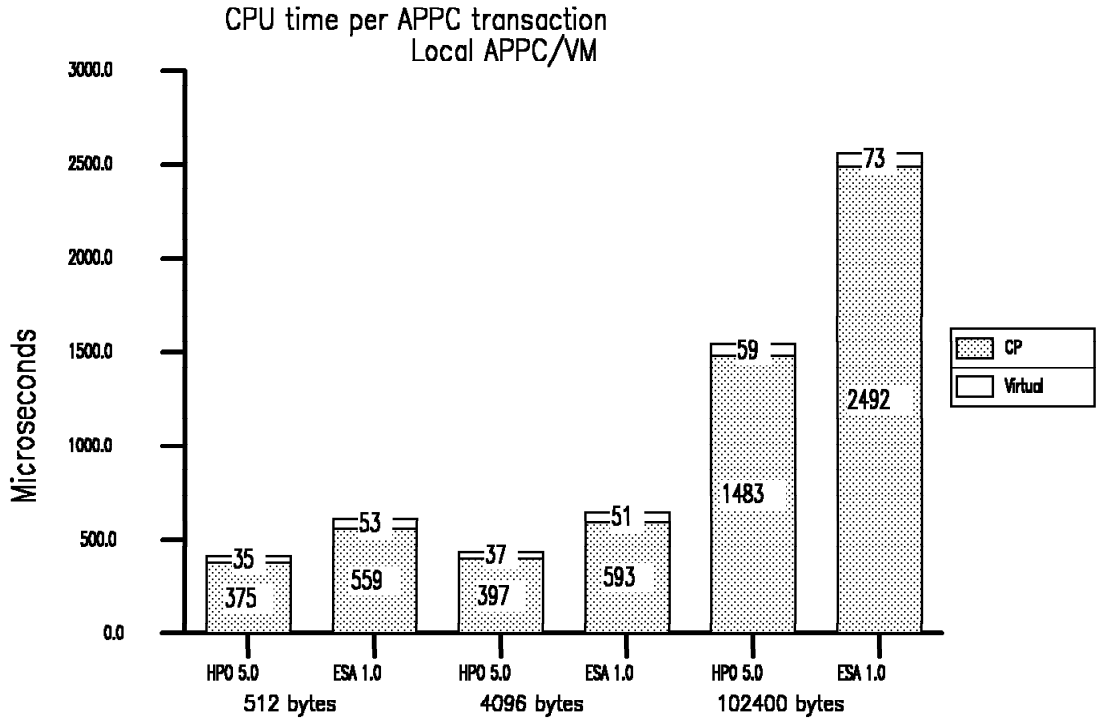

Figure 9. ESA 1.0 vs. HPO 5.0 -- APPC local with one processor: Processor busy time per command for single-thread, single processor APPC/VM measurements as described in "APPC/VM Single-user" on page 59.

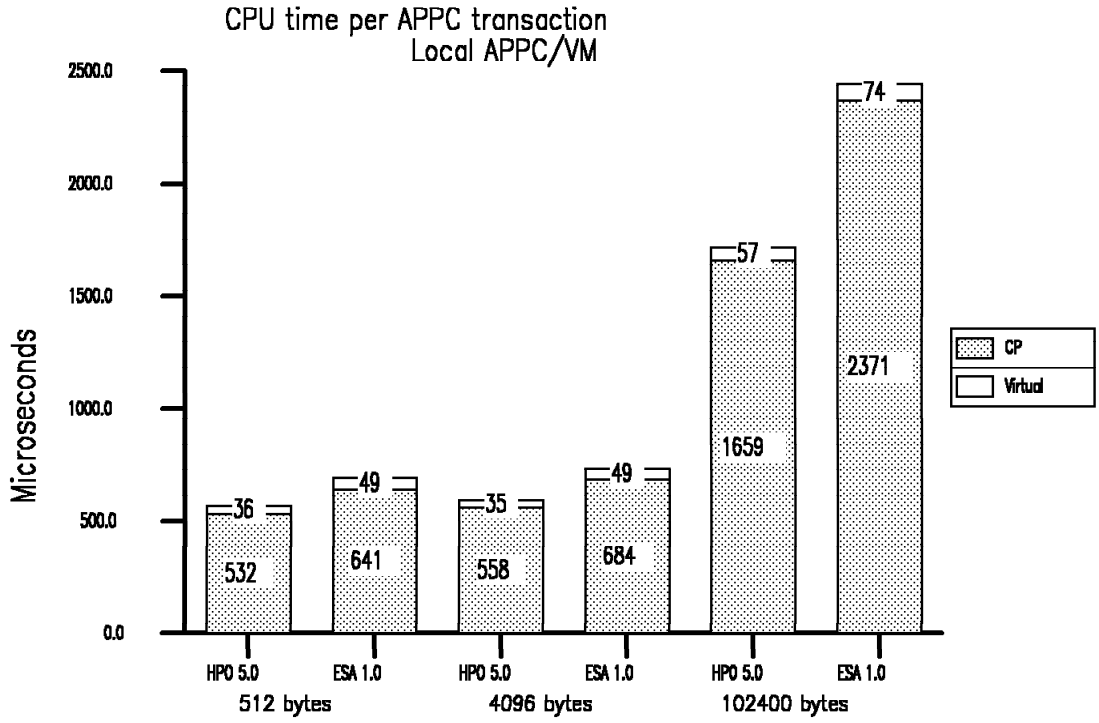

Figure 10. ESA 1.0 vs. HPO 5.0 -- APPC local with two processors: Processor busy time per command for single-thread, two processor APPC/VM measurements as described in "APPC/VM Single-user" on page 59.

Figure 11 compares VM ESA 1.0 and HPO 5.0 for two multi-user benchmarks. The benchmarks are:

- 2 resources, 150 users, 512 byte transfers;
- 2 resources, 30 users, 102400 byte transfers.

These measurements are also described in detail in "APPC/VM Multi-user" on page 62.

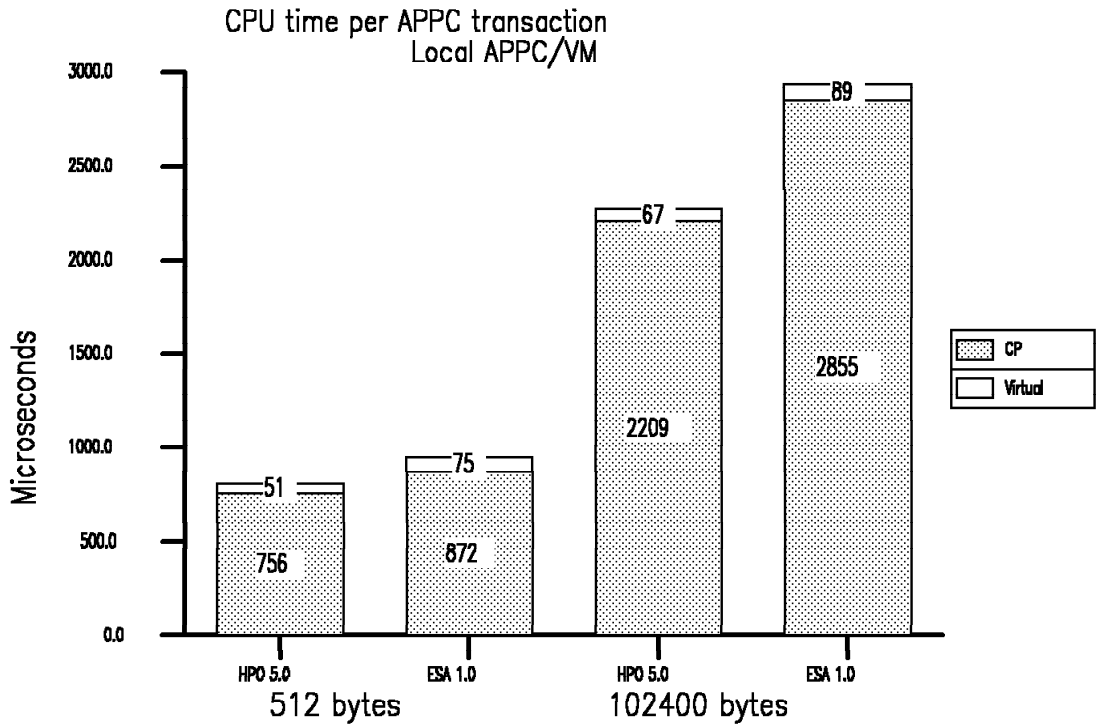

Figure 11. ESA 1.0 vs. HPO 5.0 -- APPC local Multi-user runs: Processor busy time per command for multi-user APPC/VM measurements as described in "APPC/VM Multi-user" on page 62.

Although VM ESA 1.0 is slower than HPO 5.0 in these benchmarks, it is worth pointing out some positive aspects of the measurements. In the multi-user runs, both the number of users and the transaction rate for those users was dictated by the existence of the 16M line in HPO. The first measurements were made with 64M of storage. It quickly became clear that HPO 5.0 looked artificially bad because of page thrashing across the 16M line. This thrashing obscured the measurement of APPC performance, so the test case was arranged to minimize movement of pages across the 16M line. Movement of pages across the 16M line for the measurements reported here occurs at rates of under 1 page per second.

To match the HPO configuration, the VM ESA 1.0 system was generated to use only 16M of real storage. In essence, the ESA 1.0 system was hobbled to insure that it would not profit from HPO's 16M limit. In any real environment, of course, ESA 1.0 would be able to take advantage of storage above 16M.

2.1.5.2 AVS

The AVS runs were made running VTAM 3.3 and AVS together in an XA-mode virtual machine. The environment measured included:

- 2000 users on a 3090-600J;
- 6 resources on a 3090-300S;

Our 3090-600J was able to support these 2000 users each receiving 512 bytes about 18 times per minute. The total message rate was about 606 512-byte transfers per second. Processor utilization averaged 34.3% on the 600J and 74.7% on the 300S. The CPU utilization of the VTAM/AVS virtual machines to accomplish this throughput was 206% on the 600J and 224% on the 300S. Please refer to “AVS Measurements” on page 64 for details.

2.1.6 Processor Capacity

The processor capacity study is a set of measurements of VM/ESA 1.0 on different 3090J processors. The study was on the 180J, 200J, 300J and 600J. The way the measurements were obtained was by taking a 3090-600J and varying off the number of processors needed to have the 180, 200, 300 and 600 environments. As the number of processors was increased, the number of users was increased in order to keep the average processor utilization at 90%.

The measurements that were done were in the CMS intensive environment and the users were run in 370 mode. The following figures summarize these results and section “CMS Intensive Measurements” on page 66 contains the details of the measurements.

> Examination of the data shows that performance of the measurement on the 3090-600J is out of
> line with the data for the smaller processors. This artificial degradation is the result of the use of
> the SET AUTOREAD ON command in the test scripts for command synchronization by the
> workload driver (TPNS). This causes users to enter a VMREAD state after completing commands.
> When the user enters VMREAD state it’s virtual machine issues an IUCV two-way send to the
> VSCS virtual machine and is placed on the RECEIVE queue which is a FIFO queue. The user’s
> message block stays on this queue until a reply comes, causing the queue to be searched to match
> the reply with the appropriate user’s message block. Since the queue is FIFO users with longer
> think times migrate to the top of the queue, causing a search of the entire queue in most instances.
> The increase in the number of users highlights this problem on the 6-way.

> To improve performance for the case when users are in a VMREAD state HCPIUF was changed
> so that the RECEIVE queue is now LIFO instead of FIFO. This causes users with long think
> times to migrate toward the bottom while those whose requests complete in a shorter timeframe
> can be found near the top. Now the entire queue is rarely searched giving an overall improvement
> in performance. As a result performance is more in line with expectations. See section “CMS In-
> tensive Measurements” on page 66 for a comparison of ESA 1.0 at the starter set level to the same
> level plus the aforementioned fix to HCPIUF.

> Please note that the modified HCPIUF code is prototype code. Although these measurements have
> nearly 100% of the users on VSCS’s IUCV RECEIVE queue, sampled installations had only about
> 10% of their users in this queue. Since it was determined that this modification would not greatly
> benefit real users the change won’t be available in ESA 1.0 but has been rolled into ESA 1.1. Only
> those SNA installations whose users enter VMREAD state frequently will see an improvement by
> applying this fix. Installations that have only a small fraction of their users in VMREAD state will
> see no change in their performance.

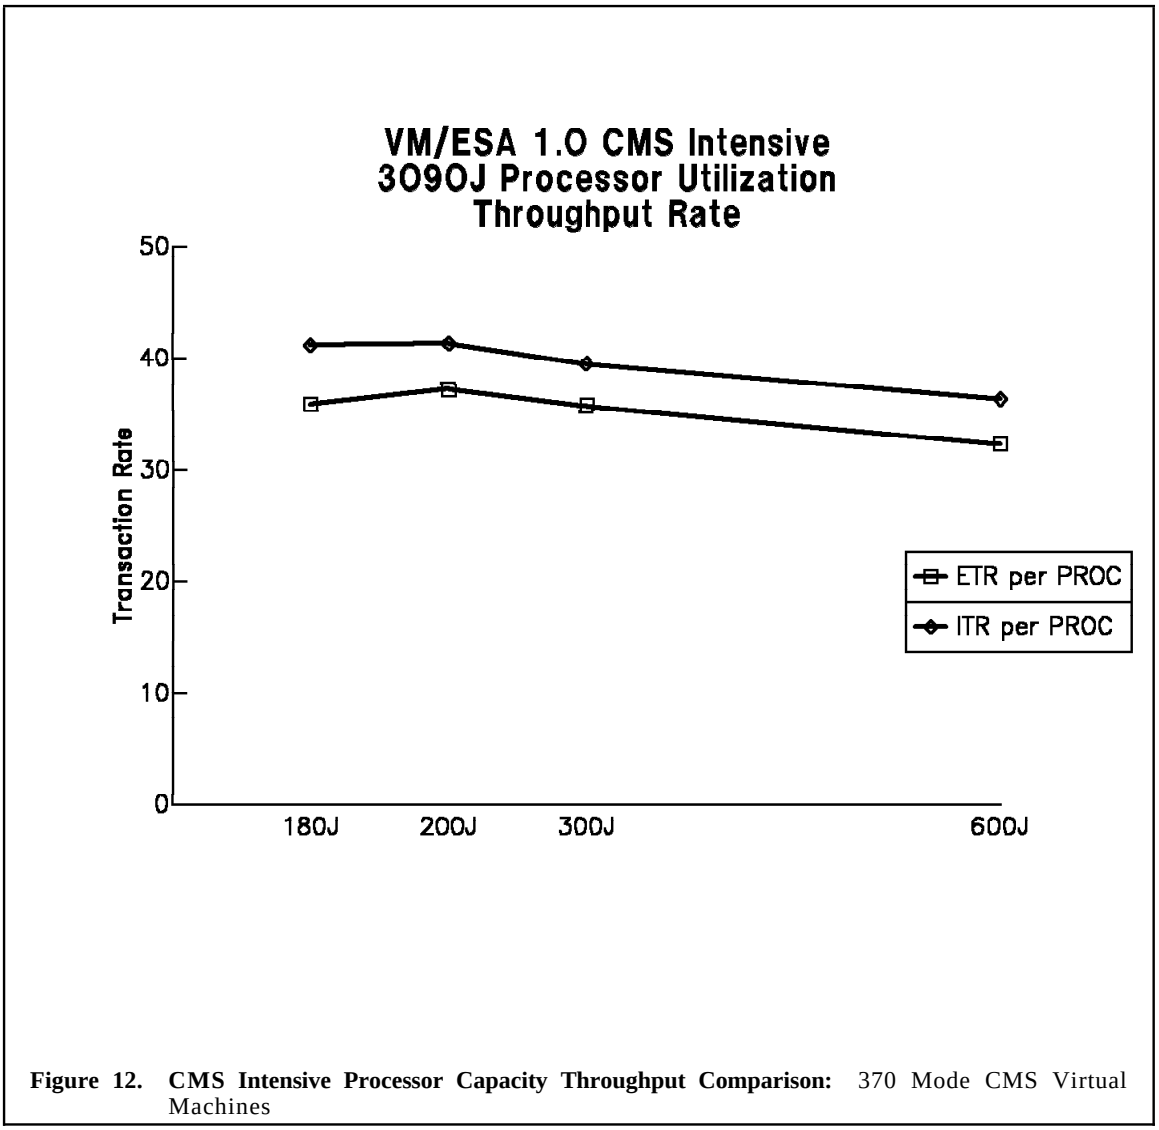

Figure 12. CMS Intensive Processor Capacity Throughput Comparison: 370 Mode CMS Virtual Machines

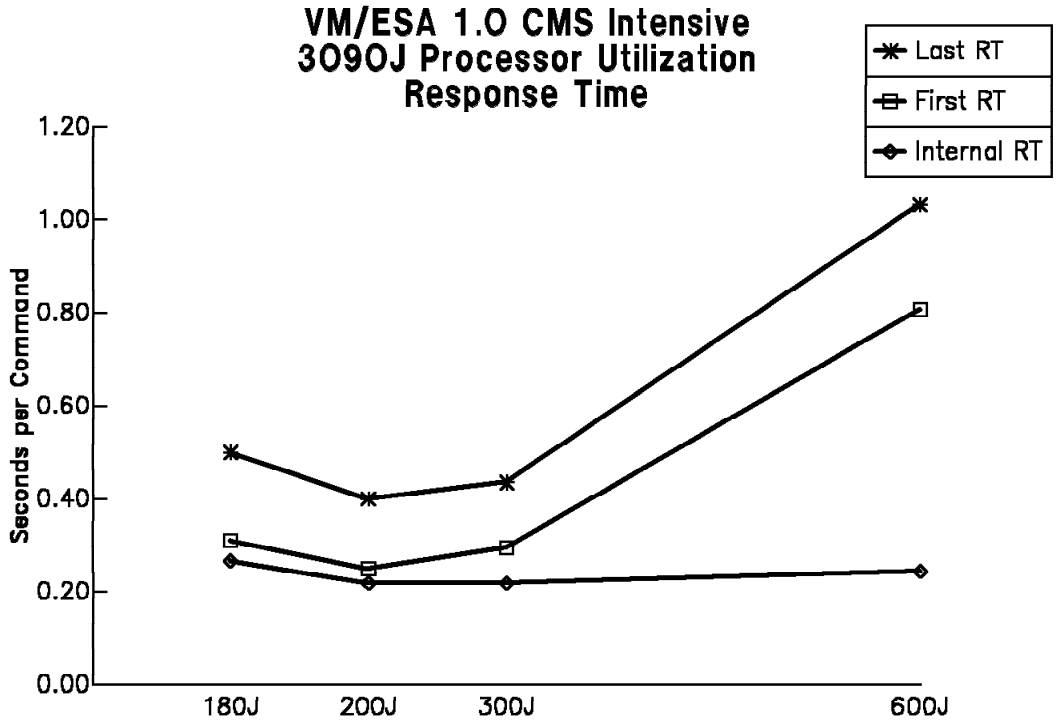

Figure 13. CMS Intensive Processor Capacity Response Time Comparison: 370 Mode CMS Virtual Machines

2.1.7 Migration

The intent of these measurements was to take the existing VM/XA/SP 2.0 and VM/XA SP 2.1 systems and migrate to VM/ESA 1.0 by changing one component at a time. This way any change in performance could be isolated to a specific component. A CMS intensive benchmark (FS7B0R) was used with the users running in XA mode.

The sequence of measurements was:

VM/XA SP 2.0 base measurement

VM/XA SP 2.1 base measurement

CP/ESA 1.0 (at a 10/89 driver level) replaced VM/XA SP CP. The measurements gave a comparison of the performance of the CPs.

GCS/XA (at a 11/89 driver level) replaced VM/XA SP GCS/370. These measurements gave comparison of the performance of GCS/XA vs. GCS/370.

CMS/ESA (at a 11/89 driver level) replaced VM/XA SP CMS 5.5. These measurements gave comparison of the performance of CMS/ESA vs. CMS 5.5.

VTAM 3.3 (at a 12/89 driver level) replaced VM/XA SP VTAM 3.2. These measurements gave comparison of the performance of VTAM 3.3. vs. VTAM 3.2.

This way migration from VM/XA SP 2.0 or from VM/XA SP 2.1 to VM/ESA 1.0 was done the way a customer would do it; changing one component at a time.

The results of these early studies show that there is an improvement in performance in VM/ESA 1.0 over either VM/XA SP 2.0 and VM/XA SP 2.1. Throughput and internal response time are better while using less processor resources.

NOTE: Since these were done on early ESA 1.0 drivers, the actual change the users may see as they migrate to ESA 1.0 may be different. Measurements with later drivers and the PID level system showed an even greater performance improvement for VM/ESA 1.0 over both VM/XA SP 2.0 and VM/XA SP 2.1. The PID level measurements were done with an increased number of users, 5400 instead of 5200. A single VTAM virtual machine was used in the later runs as opposed to 3 VTAMs in these early measurements. The PID level measurements showed additional improvements of about:

6% in internal response time and
1.5% in the internal throughput rate
while processor utilization increased by about 2%.

See the section “370 to XA Measurements” on page 47 (XA Mode) for these updated performance results obtained from those PID level measurements.

The following figures show the change in the processor utilization, the internal response time and the internal throughput rate ratio on the early drivers.

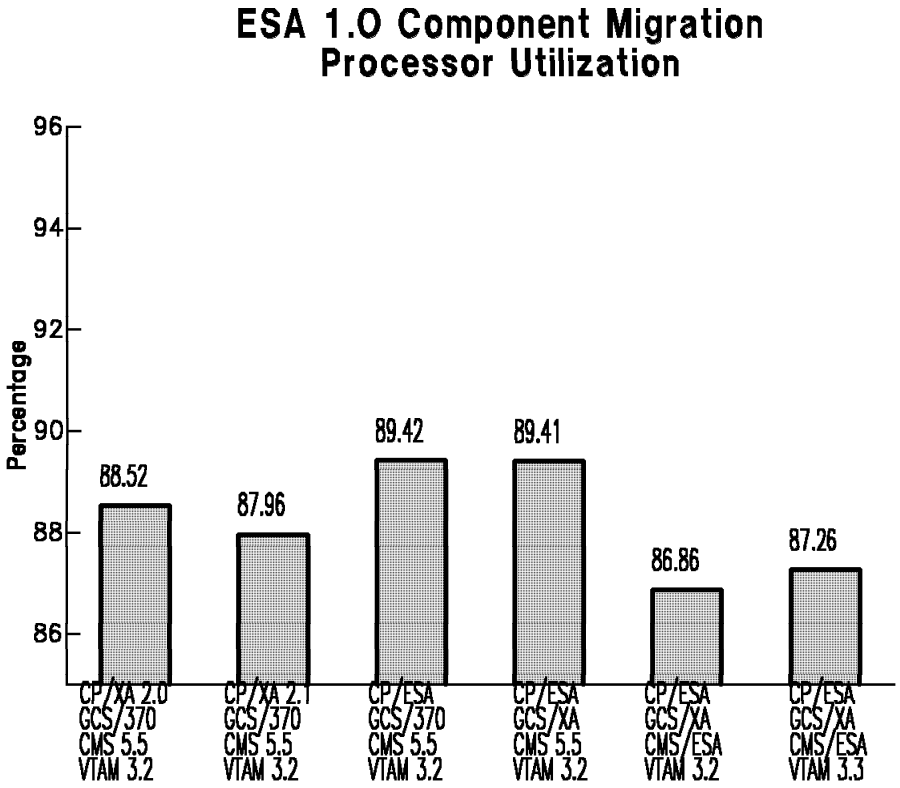

Figure 14. Migration -- Processor Utilization Comparison with 5200 Users

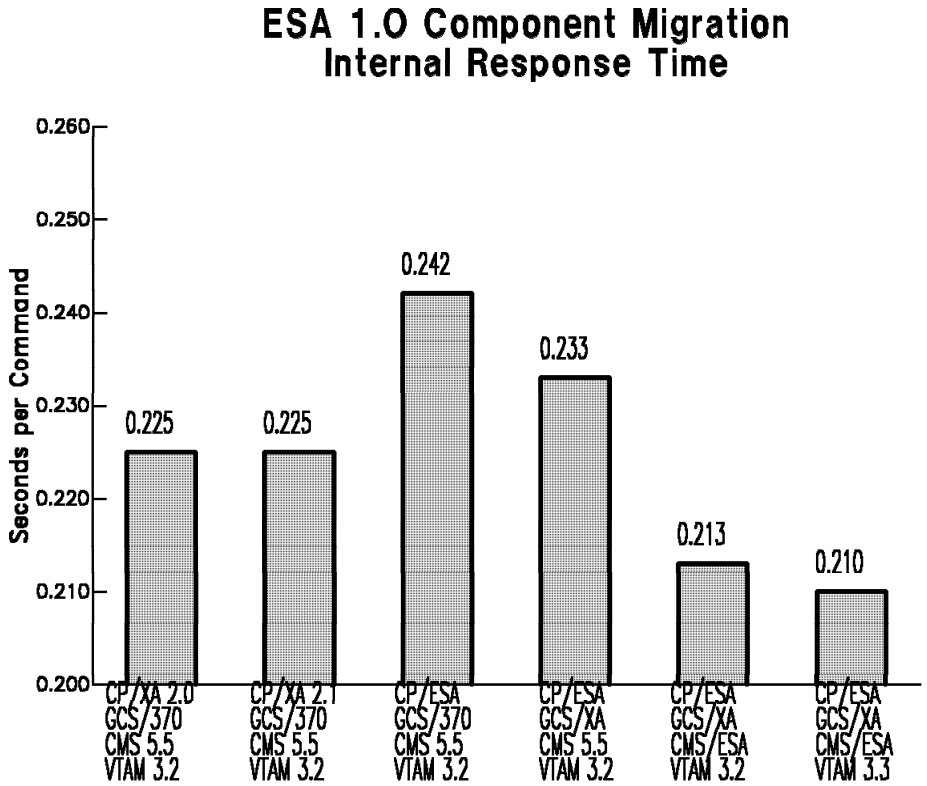

Figure 15. Migration -- Response Time Comparison with 5200 Users

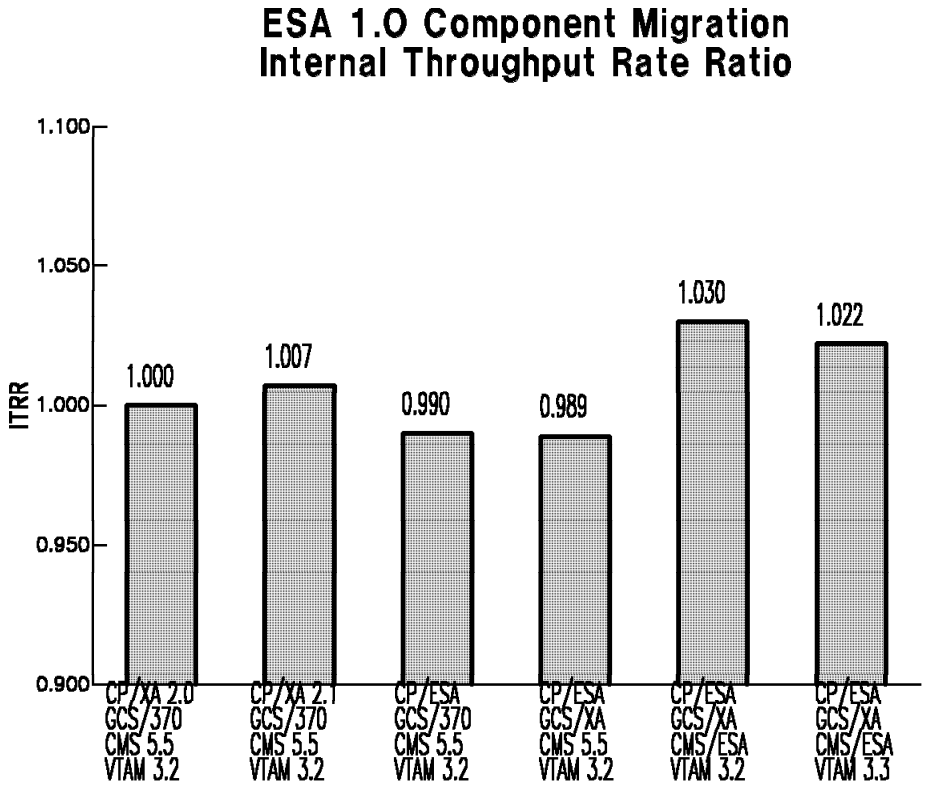

Figure 16. Migration -- ITRR Comparison with 5200 Users

2.1.8 Single VTAM

Need for single VTAM measurements: Measurements with single VTAM configurations were determined to be important in the ESA 1.0 system because of the addition of 31-bit addressing. This provides storage constraint relief in the Single VTAM environment. By examining the results of the Migration measurements in "Migration" on page 21, it can be seen that VTAM 3.2 and VTAM 3.3 were close in performance characteristics when run in a multiple VTAM configuration. This was the expected result since no changes that would affect performance were added to VTAM 3.3. However, the Migration measurements were made to compare VTAM 3.3 to VTAM 3.2, and they did not exploit the most important change which was added to VTAM 3.3. This change was the ability to use 31 bit addressing, and measuring this change was the reason that the single VTAM measurements were taken.

Requirements from users of VTAM had indicated a need for the storage constraint relief that would be provided by 31 bit addressing. It was expected that this change would allow support of more users through a single VTAM image in VTAM 3.3. Single VTAM measurements were therefore used to establish whether or not VTAM 3.3 could support more users than VTAM 3.2 from a storage point of view.

Measurements concentrated on two types of environments: first, where VTAM and VSCS shared a single virtual machine, and second, where two copies of VSCS were in separate virtual machines from VTAM. The key factor in both environments was the use of a single copy of VTAM, so that an indication of the effect of 31 bit addressing support in VTAM could be measured.

The CMS intensive workload (FS7B0R) was used for the measurements. While the results achieved using this workload indicate a substantial increase in the number of users that can be supported by VTAM 3.3 over VTAM 3.2, these numbers should not be used as absolute values for any other system or workload. For instance, when considering VTAM 3.3, the number of users that could be supported by a single VTAM could increase if the customer's workload is not as CPU intensive as the FS7B0R workload.

Single VTAM/VSCS virtual machine measurements: The purpose of the measurements on the single VTAM/VSCS virtual machine environment was to establish an upper limit on the number of users that could be supported. Storage measurements for VTAM 3.2 had previously indicated that, due to storage constraints, a maximum of approximately 1600 users could be supported by VTAM 3.2 before moving to an alternate configuration (i.e., either moving VSCS to a separate virtual machine or using multiple copies of VTAM). Results detailed in "Single VTAM" on page 73 show that over 4200 users can be supported in the same configuration by VTAM 3.3 with the FS7B0R workload.

Single VTAM and separate VSCS measurements: There were two purposes for the measurements done in the configuration which contained a single VTAM and two separate VSCS virtual machines. First, a comparison of the single VTAM to a multiple VTAM configuration running the same number of users was needed. Details in "Single VTAM" on page 73 show that many key characteristics of the two configurations were very close in measurements made with 5700 users. There was some degradation in response time with the single VTAM configuration, however. Second, the purpose for the single VTAM with two separate VSCS virtual machine configuration measurements was to determine how many users could be supported by this environment using the FS7B0R workload. A measurement with 7760 users was completed for this purpose. This should not be considered an upper limit for this configuration, as the measurement was constrained by the network configuration. Addition of 3745s to replace the 3725s used should increase the number of users that could be supported somewhat.

The configuration which included one VTAM and two separate VSCS virtual machines was chosen for a couple of reasons. First, this configuration could take advantage of the multiple processors in the 3090-600J machine. This is due to the fact that the VSCS virtual machines could potentially be scheduled on separate processors at the same time as the VTAM virtual machine was scheduled, thus overlapping processing time. Second, more than two separate VSCS virtual machines were not used in any of the measurements. This was because measurement results showed that CPU

resources were not used to the full capacity with the two VSCSs (and thus a third one was not required for the measurements). **NOTE:** VTAM capacity itself was not a limiting factor in any of the VTAM capacity measurements. Limitations of the run environment prevented the determination of VTAM's limit.

2.1.9 MVS Guest

Guest operating system performance on VM/ESA 1.0 was predicted to be equivalent to VM/XA SP 2.1 because all of the guest performance improvements made to VM/XA SP 2.0 and VM/XA SP 2.1 were propagated into VM/ESA 1.0. These included numerous SIE, RRF, PR/SM and MPG enhancements.

The capacity of MVS V=R, V=F, and V=V guests in a local environment running under VM/ESA 1.0 on a 3090-180J processor was measured using the CB84 workload. See "MVS Guest (CB84)" on page 180 for CB84 workload description. Since no performance gains were expected, these measurements can be viewed as regression measurements using VM/XA SP 2.1 as the comparison base. The results of the CB84 measurements verified that guest performance had not degraded from that of either VM/XA SP 2.0 or VM/XA SP 2.1.

The following figure is a graphic summary of the ITRR comparisons between MVS guests, running under the three different systems (VM/XA SP 2.0, VM/XA SP 2.1, VM/ESA 1.0), and MVS running in a native environment. While some slight differences in the ITRRs are evident, these are quite small and are not significant. The conclusion to be drawn is that the performance gains experienced by MVS guests running under either VM/XA SP 2.0 or VM/XA SP 2.1 have been carried forward into VM/ESA 1.0.

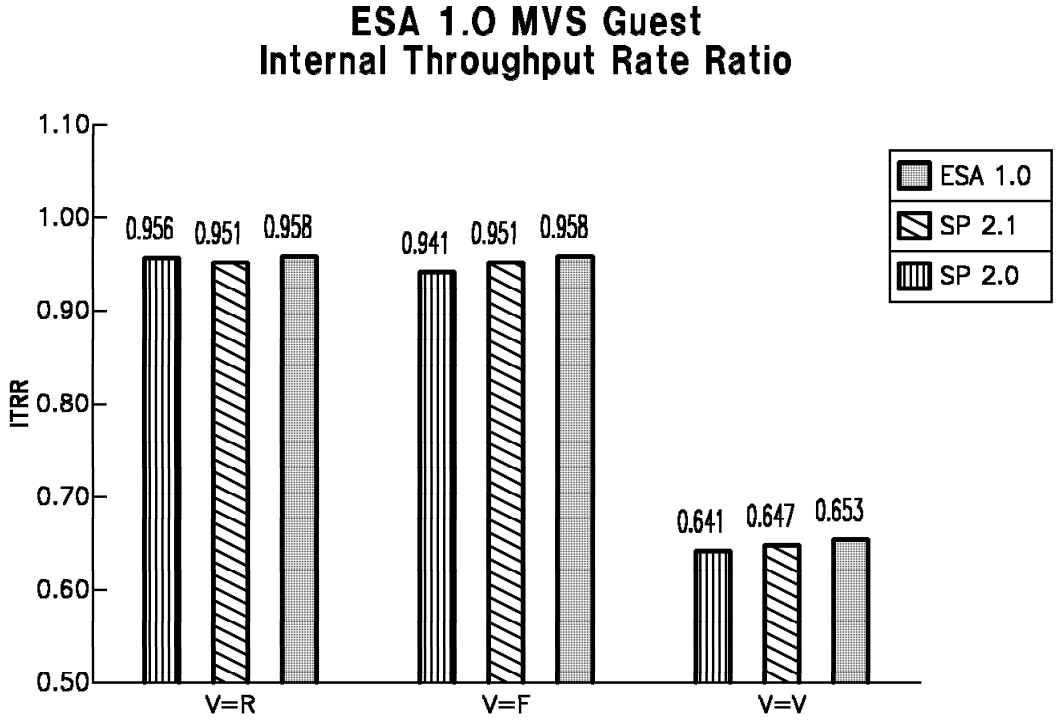

Figure 17. MVS Guest ITRR Comparisons

2.2 VM/ESA 1.0 370 Feature

2.2.1 CMS Regression

The purpose of this section is to summarize performance differences that have occurred due to the addition of new function and service since the introduction of VM/SP Release 6. The workload executed was not changed explicitly to use any of the new ESA 1.0 functions. Therefore the same end user work was completed in both cases. These comparisons can be used as an indicator of what might be expected with similar workloads when up-leveling to this new level of VM. Comparisons were made with varying usage of the Shared File System on various processors and storage sizes.

The storage required by CMS has changed in the following ways: The CMS shared nucleus area has grown, the non-shared low nucleus area has grown, new CRR control blocks were added, and the storage referenced for moderate to large accessed disks has been reduced. The net of these changes is very workload dependent. The workload and measured environments showed paging rate changes ranging from +11% to -23%.

The Processor Busy Time (PBT) used by CMS has also changed. Several individual changes were made, with the most notable ones being reduced processor busy time for EXECIO when called from REXX, COPYFILE improvements on a per record basis, Assembler time reductions due to an improvement in the CMSSTOR macro expansion, and an increase in the IPL CMS time. The EXECIO improvement is also responsible for the large reduction in the PRIVOP rate. The SPKA instructions were eliminated from the linkage path between the REXX and EXECIO modules. For those measurements using the Shared File System, a reduction in file pool requests was seen due to the elimination of most of the user virtual machine requests to update the FST information for accessed directories. The environments measured showed processor busy time reductions ranging from -1% to -8%.

The virtual file I/O rate showed a slight reduction (-3% to -7%) mainly due to a reduction in I/Os for the Assembler when expanding the CMSSTOR macro.

The combination of the above performance changes netted a change in average response time of +4% to -20%.

The data in each section

″4381-13 / 16M / 0% SFS″ on page 84
″4381-13 / 16M / 35% SFS″ on page 88
″4381-13 / 16M / MAX% SFS″ on page 91

shows what happens when moving to the new release for a given SFS usage (0%, 35%, or MAX%). By comparing across the three sections, the data shows that the degree of performance improvement decreases with increasing SFS usage. This trend is summarized in the following table:

Average Percentage Change
Relative to SP6

SFS usage:	0%	35%	MAX%
processor:	-7	-5	-3
storage:	-7	0	+6
file I/Os:	-6	-5	-4

A measurement was also run with CMS interactive users in 768K of real storage. This was done with the small CP option. The high paging rate on a per command basis indicated a severely storage constrained environment. However, the resulting response times were acceptable.

2.2.2 EDF and SFS Comparisons

Measurements were obtained on ESA 1.0 370 Feature to compare the 0% SFS (EDF) workload, the 35% SFS workload (all user files are in SFS), and the MAX% SFS workload (all files that SFS supports are in SFS). For a more complete description of these workloads, see "Appendix C. Workloads" on page 163.

Two views are provided:

1. Equal Utilization: EDF and SFS measurements are compared at the same processor utilization. At 35% SFS, this view showed the following:

   - A 13% average reduction in the number of users was required to hold CPU utilization constant.
   - Processor busy time per command increased an average of 18%.
   - Paging per command increased an average of 13%.
   - Virtual file I/Os per command increased an average of 16%.
   - The total response time per command increased from an average of 0.8 seconds to an average of 0.9 seconds.

2. Equal Users: EDF and SFS measurements are compared at the same number of users. This view showed:

   - Processor busy time per command increased an average of 24%.
   - Paging per command increased by 44% to 60%.
   - Virtual file I/Os per command increased an average of 17%.
   - Total response time per command increased from an average of 0.8 seconds to an average of 1.4 seconds.

The measurements show that SFS requires more processor time and real storage than EDF when performing similar tasks. However, SFS has significant functional improvements over minidisk and SFS can provide significant DASD savings when compared to minidisk (for more details, see "EDF and SFS Comparisons" on page 13).

Notice that response times for EDF and SFS are similar with similar processor utilization.

Note: A Coordinated Resource Recovery (CRR) server did exist for these measurements, but in this environment the recovery server is not involved in mainline processing. If the recovery server had not been running, limp mode could increase processor requirements by as high as 40%. To have acceptable performance you must have a CRR server running, even in regression environments where you are not using CRR.

2.2.3 Connectivity

All three connectivity environments listed in "Connectivity" on page 14 were measured using the Connectivity Driver described in "Connectivity Driver Tool" on page 176. These measurements were all multi-user runs and were performed to insure that there was little, if any, degradation from SP 6.

The slight degradation that does occur has two sources:

1. a small increase in CP IUCV pathlength due to new APPC functions; and

2. a small increase in CP pathlength due to >16M support.

2.2.3.1 APPC/VM

The Connectivity Driver was used to create multi-user loads running transactions of 512, 4096, and 102400 bytes. The smaller transaction sizes showed a modest decrease in processor time, and the 102400-byte case showed a slight increase. Figure 18 shows the results of local APPC/VM measurements made on a 4381-13.

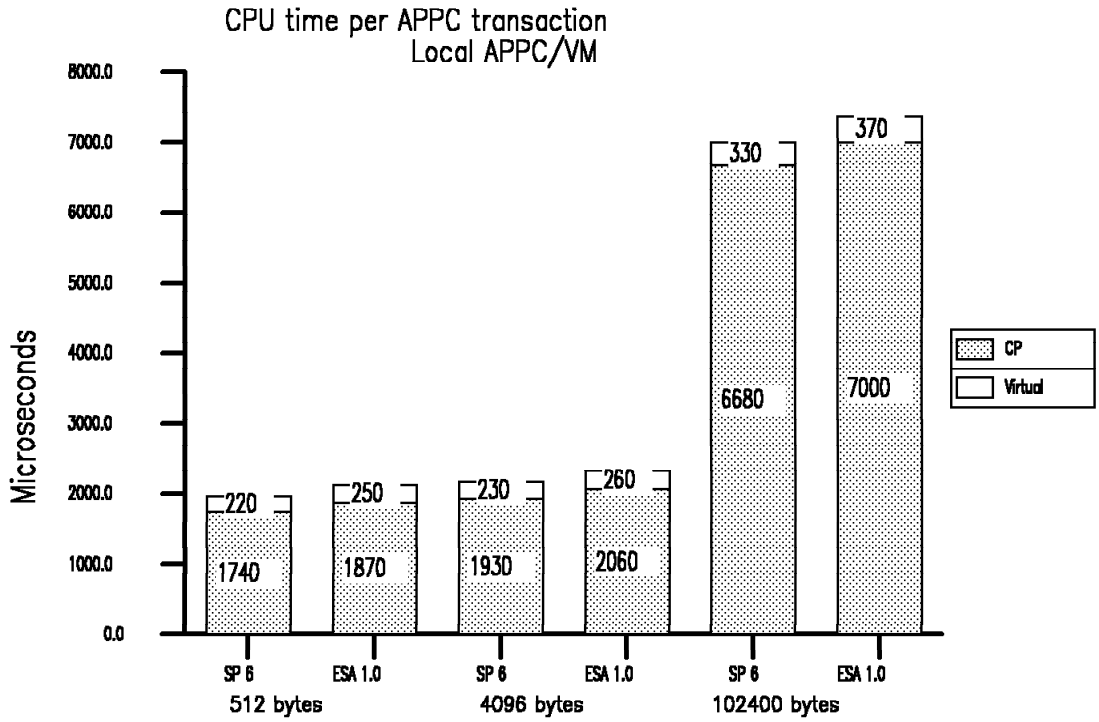

Figure 18. ESA 1.0 370 Feature vs. SP 6 -- APPC local: Multi-thread measurements as described in "APPC/VM Measurements" on page 107.

2.2.3.2 AVS

AVS was used to connect the users and resources on two different systems. Figure 19 shows the small increase in processor busy time per transaction when using AVS. Both systems - user and resource - are included in the data represented by this figure. The resource side runs on a 4381-13 and the user side runs on one processor of a 4381-92E.

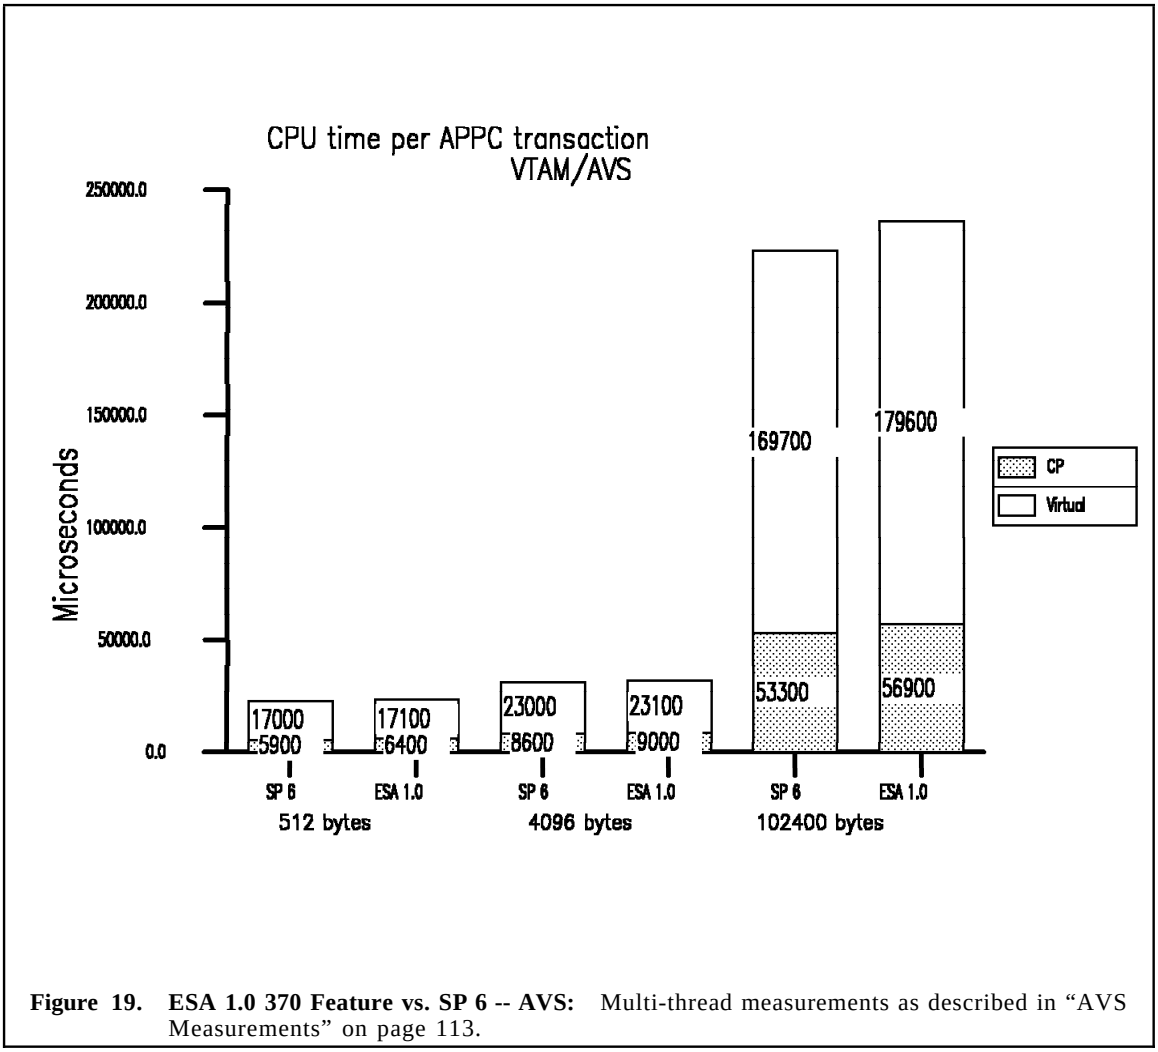

Figure 19. ESA 1.0 370 Feature vs. SP 6 -- AVS: Multi-thread measurements as described in "AVS Measurements" on page 113.

2.2.3.3 TSAF

TSAF was used to connect users and resources on two systems in a TSAF collection. Figure 20 compares TSAF for ESA 1.0 370 Feature and SP 6. Both systems - user and resource - are included in the data represented by this figure. The resource side runs on a 4381-13 and the user side runs on one processor of a 4381-92E. Improvements to TSAF cause a reduction in processor time per command, particularly in the case of large transaction sizes (102400 bytes).

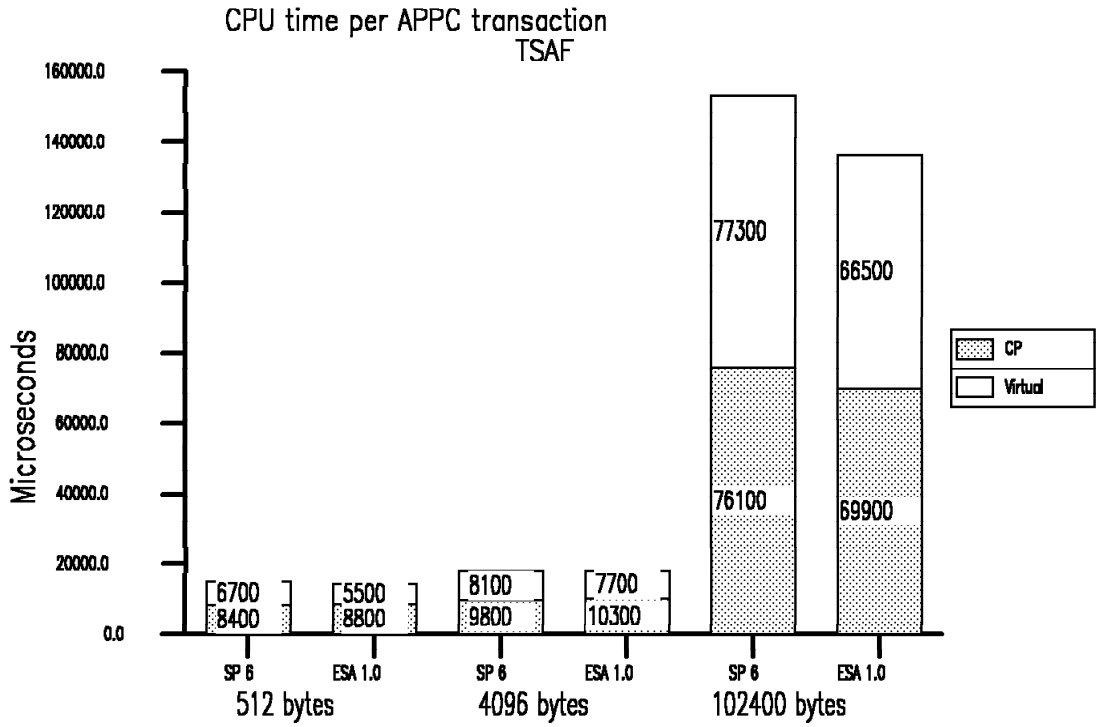

Figure 20. ESA 1.0 370 Feature vs. SP 6 -- TSAF: Multi-thread measurements as described in "TSAF Measurements" on page 119.

2.2.4 Coordinated Resource Recovery

Coordinated Resource Recovery (CRR) provides support in the CMS environment to update multiple resources and have all updates committed or rolled back together. It implements the SNA LU 6.2 Syncpoint Architecture. CRR function is available on all VM/ESA 1.0 systems. It is discussed here in the "VM/ESA 1.0 370 Feature" section only because the measurements and analysis were made on a ESA 1.0 370 Feature system. Observations made in this section also apply to ESA 1.0.

Part of CRR runs in a server virtual machine, referred to as the Recovery Server. A piece of CRR also runs in the end user machine in both the CMS nucleus and Callable Services Library (CSL) code, and is called the syncpoint manager (SPM).

In order for resources to be coordinated by CRR, they need a resource adapter which runs in the end user virtual machine. A resource adapter is the interface between the resource manager and the SPM. The resource manager is the facility which controls the resource being coordinated. For SFS, this would be the SFS file pool server. VM/ESA 1.0 provides two resource adapters: an SFS adapter and an APPC adapter. The APPC adapter is known as the Protected Conversation Adapter (PCA). Support for protected conversations is provided by new function in the CP and AVS components of VM/ESA 1.0.

Depending on the scenario, the CRR processing requirements vary. There are three levels of coordination in CRR:

1. Coupled Commit - this requires internal interfaces designed for SFS exploitation. The SPM is notified prior to a coupled commit by the resource adapter. Under certain conditions, the SPM will allow SFS to couple the commit request on another SFS file pool server request. When the request is complete, the adapter notifies the SPM of the results.
2. Simple or One-phase Commit - the SPM coordinates the commit by driving participating resources just once. In regression environments, the CSL Commit function will result in a Simple Commit.
3. Two-phase Commit - the SPM coordinates commit by driving participating resources through two phases. One phase polls all resources on their ability to commit and the second phase causes the commit processing.

Regression Performance: CRR impacts regression performance when comparing VM/SP Release 6 to VM/ESA 1.0. For an all minidisk environment, the impact is negligible. For an SFS environment, CRR is approximately 5% of the virtual pathlength for the average command. In regression environments (no coordinated commits), the Recovery Server is not involved in mainline processing. All CRR processing is performed in the end user virtual machine by the SPM.

One important aspect of CRR in regression environments is the need for a CRR Recovery Server. *Limp Mode* is when the Recovery Server is unavailable. In this state, regression coordination (coupled commit or simple commit) processing can continue, but not two-phase coordination. There is a performance degradation associated with Limp Mode. This is due to the high overhead of determining if a Recovery Server exists, which is repeated for each CMS command that involves coordination. The Limp Mode degradation can be as high as a 40% increase in CPU requirements.

New Function Performance: CRR new function performance concerns environments where simple and two-phase commits are being processed. Use of CRR new function is expected to be low in most environments. However, even with moderate usage, continual performance monitoring and tuning should not be required, unless recommendations are not followed. Occasional monitoring is recommended to establish base information and determine the level of CRR usage.

Performance considerations for environments with high CRR usage are discussed here. The following are factors of performance impact:

- Type of syncpoint (simple or two-phase)
- Number/type of resources
- Syncpoint rate
- Log disk size/placement

The type of syncpoint processing is a key factor in the performance of systems using CRR. Two-phase commits are more expensive than simple commits. The difference depends on other performance factors discussed later in this section, but two-phase commits can be an order of magnitude more expensive. This increase is due to 1) the involvement of the CRR Recovery Server for logging in two-phase and 2) the additional communication required between the SPM and participating resources. A simple commit can be used if the following are all true:

1. No protected conversations participating.
2. One or less resources have been updated since last syncpoint.
3. All participating resources support simple commit.

Besides determining whether a simple or two-phase commit is required, the number and type of participating resources can affect performance. In general, overhead increases as the number of

resources increases. Resources that have been updated result in higher overhead compared to read-only resources. Note that protected conversations are always considered updated resources and result in two-phase commits.

The rate of commits or syncpoints may impact performance. There are a number of optimizations for log processing in the Recovery Server. These optimizations may result in better performance as the syncpoint rate increases. Figure 21 on page 34 illustrates this effect. Concurrency of commit processing is important. A system with a single user who issues commits at a higher rate will not see an improvement since the opportunities for overlap do not exist.

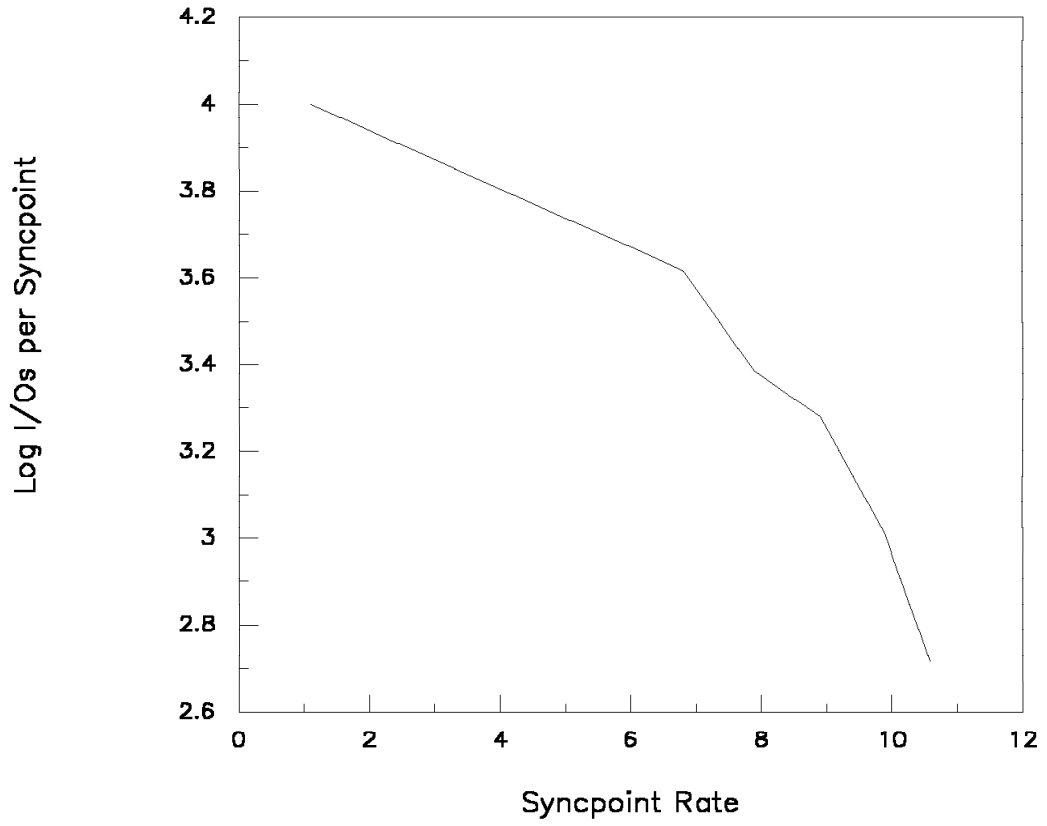

Figure 21. Log I/Os per Syncpoint vs. Syncpoint Rate: This graph illustrates the logging efficiencies that can be realized when processing rate increases.

The CRR Recovery Server uses dual logging for reliability reasons. The placement of these log disks is crucial for good performance. They should be placed on separate channels and control units if possible. Since the majority of I/O activity to them is write, they do not gain any advantage from cache DASD or minidisk caching.

2.2.5 GCS

This section describes measurements comparing ESA 1.0 370 feature running GCS6 to ESA 1.0 370 feature running GCS370. The purpose of making these measurements was simply to insure that performance of GCS had not degraded significantly from the previous release.

The results show GCS370 to be equivalent to GCS6.

2.2.6 Greater Than 16M Real Storage

The ESA 1.0 370 Feature provides storage and channel constraint relief to the VM/SP customer. Configurations that are currently storage constrained with 16M will now be able to access up to 64M of real storage, similar to VM/HPO. This extra storage can substantially reduce the paging rate as well as reduce response times. The extended channel support provides channel constraint relief which allows greater availability to support more channels.

All ESA 1.0 370 Feature measurements for extended storage support were conducted on a 4381-13 with 32M of real storage. These results were compared to SP 6 with 16M of real storage and to ESA 1.0 370 Feature with 16M of real storage. Three intervals were run at 70%, 80%, and 90% CPU utilization with the FS7B35 workload.

As expected, the new support reduces the average response time because of a reduction in the number of page I/Os. The average response time decreased 39% to 53% depending on the level of system loading when compared to SP 6 and 34% to 46% when compared to ESA 1.0 370 Feature.

Overall processor busy time per command decreased. This is due to a large savings in page fault handling, which results in lower CP CPU consumption. There is a slight increase in supervisor CPU time for the movement of pages across the 16M boundary line and to check the location of pages. PBT/CMD decreased 13% to 17% when compared to SP 6 and 8% to 12% when compared to ESA 1.0 370 Feature.

Since the above measurements were compared to only 32M of real storage, some further improvements would be observed if 64M of real storage were used.

3.0 Specific Measurements: VM/ESA 1.0

3.1 Introduction

This chapter contains the configuration details and specific results for those measurements obtained to evaluate the performance of VM/ESA 1.0 on large systems.

3.1.1 Format Description

For each group of measurements in this section there are two parts. The first part describes the run environment and the second part describes the data. A description of the four sections that make up the run environment are:

1. Workload: Gives the name of the workload associated with the data. For more detail see section “Appendix C. Workloads” on page 163.
2. Hardware Configuration: This section gives a high-level overview of the hardware configuration. It contains the following subsections:
   - Processor: The processor for which the data was collected.
   - Storage: The amount of real and expanded storage available on the processor.
   - System DASD: The packs on which the system resides.
   - Page DASD: The number of paging packs on the system and the number of channels they are spread across.
   - Spool DASD: The number of spool packs on the system and the number of channels they are spread across.
   - TDISK DASD: The number of packs containing temporary disk space on the system and the number of channels they are spread across.
   - Swap DASD: The number of swap packs on the system and the number of channels they are spread across.
   - User DASD: The number of user packs on the system and the number of channels they are spread across.
3. Software Configuration: This section contains pertinent software information. It contains the following subsections:
   - Driver: The tool used to run the workload.
   - User Think Time: Average amount of time between user interactions. Most measurements use a 30 second bactrian think time distribution. The bactrian think time distribution represents a combination of both active and inactive user think times. The distribution includes those long think times that occur when the user is not actively issuing commands. Actual user data was collected and used as input to the creation of the bactrian distribution.

The think time was reduced to 15 seconds for the HPO measurements. This was done because with a 30 second think time the HPO system ran out of real storage before the processor utilization target of 90% was obtained. With a 15 second think time, the desired 90% utilization was obtained.

The connectivity measurements used a 5 second random think time distribution to achieve the desired processor utilization.

- CMS Blocksize: The blocksize of CMS minidisks.
- User VM Size: The storage size of the user virtual machine.
- User RELSHARE: The relative share of the system resources to be scheduled for the user's virtual machine.
- Server Machines: Lists the type of server machine, how many there are, their storage size, scheduling relative share, and any special options used.

4. Measurement Summary: This contains a brief overview of the measurement data.

Refer to the glossary for definitions of performance terms used in this chapter.

3.1.2 Tools Description

A variety of tools were used in executing and evaluating the performance measurements. Some were required to validate the performance criteria. Others were used to gain a better understanding of system interactions and to validate system behavior.

The tools that were used to gather the data for this section are:

- For the HPO measurements: VMMAP, an internal use only hardware monitor, TPNS and a TPNS Log Tape Reduction Program.
- For the CMS intensive measurements on the VM/XA systems: VM/XA System Monitor, RTM/SF, an internal use only hardware monitor, TPNS, TPNS Log Tape Reduction Program, and VMPRF. For measurements that included the Shared File System, the QUERY FILEPOOL STATUS command was used as well.
- For the MVS Guest measurements: an internal use only hardware monitor, RTM, and RMF 4.1.

VMMAP reduces VM/370 MONITOR data. VMMAP is a program product which produces reports summarizing the performance and utilization of the system. The hardware monitor is an internal use only tool that collects branch, event and timing data. TPNS (Teleprocessing Network Simulator) is a terminal and network simulation tool. The TPNS Log Tape Reduction Program provides performance, load and response time information. RTM (Real Time Monitor) provides on-line performance analysis and determination facilities for VM systems. VMPRF (VM Performance Reporting Facility) is the VM/XA Monitor reduction program.

3.2 Regression

> 3.2.1 Introduction

> Regression measurements were done for the following migration paths:

> VM/HPO 5.0 to VM/ESA 1.0

> VM/XA SP 2.0 to VM/ESA 1.0

> VM/XA SP 2.1 to VM/ESA 1.0

> The CMS intensive environment was used for these measurements. The details of these measure-
> ments follow.

> 3.2.2 CMS Intensive Measurements

> *3.2.2.1 3090-200J / 64 and 256M / VM/HPO 5.0*

> The measurements on the 3090-200J comparing VM/HPO 5.0 to VM/ESA 1.0 were done in the
> following manner: First, a measurement was done to determine how many users, running the CMS
> Intensive workload, it took to run VM/HPO 5.0 at about a 90% processor utilization on the
> 3090-200J. This same number of users (1140) and the same think time (15 seconds) were then used
> for the VM/ESA 1.0 measurement.

> **NOTE:** The think time had to be reduced to 15 seconds. With a 30 second think time (used for
> the rest of the CMS Intensive measurements on VM/ESA 1.0) HPO 5.0 becomes real storage
> constrained.

> Following is the description of the environment used to do the 3090-200J CMS Intensive VM/HPO
> 5.0 vs. VM/ESA 1.0 measurement.

```
> 1) WORKLOAD:     FS7B0R

> 2) HARDWARE CONFIGURATION

>    - PROCESSOR:  3090-600J   CPU 0,1 only

>    - STORAGE:
>       - HPO:
>         - RSTOR:   64M
>         - XSTOR:   512M
>       - ESA:
>         - RSTOR:   64M and 256M
>         - XSTOR:   512M

>    - DASD:
>      - SYSTEM    PACK NAME  TYPE  CHANNEL
>                  ---------  ----  -------
>        - HPO:    PERF21     3380       5
>                  PERF2C     3380       5
>                  WKLD01     3380      16
>                  WKLD02     3380      18
>        - ESA:    PSYS02     3380       5
>                  WKLD01     3380      16
>                  WKLD02     3380      18

>                  NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
>                  ---------------  ----  ------------------
>      - PAGE:           10         3380            3
>      - SPOOL:           4         3380            2
>      - TDISK:           6         3380            2
>      - SWAP:
>        - HPO:           4         3380            2
>        - ESA:           0
>      - USER:
>        - HPO:          28         3380            7
>        - ESA:          80         3380            7

>    - COMMUNICATIONS:  CONTROLLER  NUMBER  LINES  LINESPEED
>                       ----------  ------  -----  ---------
>                         3745         1      50     56K

> 3) SOFTWARE CONFIGURATION

>    - DRIVER:           TPNS
>    - USER THINK TIME:  BACTRIAN (Avg 15 Secs)
>    - CMS BLOCK SIZE:   4K
>    - USER VM SIZE:     2M
>    - USER RELSHARE:    100

>                        TYPE      NUMBER  VM SIZE  RELSHARE  OTHER OPTIONS
>                        ----      ------  -------  --------  -------------
>    - SERVER MACHINE:   VTAM/VSCS   1       64M     10000    QUICKDSP ON

> 4) MEASUREMENT SUMMARY
```

> Two ESA 1.0 measurements were taken for this comparison. The first compares ESA 1.0 to HPO
> 5.0 with equivalent real storage set at the HPO maximum of 64M. In this case average response
> time is 10% better for ESA 1.0 because of CMS changes and MDC which yield less I/Os, but ESA
> 1.0 still costs 3% more in ITR. Also ESA 1.0 has 3754 fewer pageable pages because of the in-
> creased free storage requirements of ESA (XA), which amount to 3.3 pages per user.

> Since the 3090-200 can support up to 256M, an ESA measurement was taken to exploit the addi-
> tional real storage. When compared to HPO, average response time improved by 15% and ITR
> is only 1.0% worse. This demonstrates the growth potential of ESA over HPO since HPO tops
> out at 1200 users, but ESA could support many more users at equivalent response time.

		HPO 5.0	ESA 1.0	ESA 1.0
Storage				
REAL		64M	64M	256M
XSTOR		512M	512M	512M
Run ID		Y25R1141	Y23R1143	Y23R1141
Variables				
USERS		1140	1140	1140
VTAMs		1	1	1
VSCSs		1	1	1
PROCESSORS		2	2	2
Response Time				
TRIV INT	(VM)	0.043	0.046	0.045
NONTRIV INT	(VM)	na	0.391	0.347
TOT INT	(VM)	0.228	0.258	0.230
AVG FIRST (T)		0.330	0.300	0.280
AVG LAST (T)		0.520	0.490	0.450
Throughput				
ETR (T)		76.88	78.51	78.86
ITR (H)		43.95	42.58	43.50
ITR		na	44.19	44.95
EMUL ITR		na	68.99	68.81
ITRR (H)		1.000	1.000	0.990
ITRR		na	na	na
Proc Usage				
PBT/CMD (H)		21.478	22.646	22.262
CP/CMD (H)		8.850	8.633	8.052
EMUL/CMD (H)		12.628	14.010	14.206
Processor Util				
TOTAL (H)		174.91	184.41	181.30
UTIL/PROC (H)		87.46	92.20	90.65
TOTEMUL (H)		102.83	114.09	115.69
TVR (H)		1.70	1.62	1.57
Storage				
WKSET (V)		68	47	60
PGBLPGS		13115	9361	57459
SHRPGS		372	814	856
Paging				
READS/SEC		106	102	90
WRITES/SEC		31	24	8
PAGE/CMD		1.738	1.605	1.243
XSTOR IN/SEC		813	790	373
XSTOR OUT/SEC		823	860	405
XSTOR/CMD		20.746	21.016	9.866
I/O				
VIO RATE		1123	655	666
VIO/CMD		14.74	8.044	8.178
MDC READS		na	450	461
MDC WRITES		na	207	205
MDC MODS		na	175	178
MDC HIT RATIO		na	0.93	0.94
PRIVOPs				
PRIVOP/CMD		na	20.756	20.602
DIAG/CMD		na	16.707	16.570
DIAG 08/CMD		na	0.726	0.723
DIAG 10/CMD		na	5.464	5.186
DIAG 58/CMD		na	1.223	1.217
DIAG A4/CMD		na	3.923	3.956
DIAG A8/CMD		na	1.911	1.966
Note: T=TPNS, V=VMPRF, VM=VMMAP, H=Hardware Monitor, Unmarked=RTM				

Table 1. VM/HPO 5.0 vs. VM/ESA 1.0 on a 3090-200J

3.2.2.2 3090-600J / 512M / VM/XA SP 2.0

Following is a description of the environment used for the CMS Intensive regression measurements on the 3090-600J comparing VM/XA SP 2.0 and VM/ESA 1.0.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

- PROCESSOR: 3090-600J
- STORAGE:
  - RSTOR: 512M
  - XSTOR: 2G

- DASD:

	PACK NAME	TYPE	CHANNEL
- SYSTEM:	PSYS02	3380	5
	WKLD01	3380	16
	WKLD02	3380	18

	NUMBER OF PACKS	TYPE	NUMBER OF CHANNELS
- PAGE:	20	3380	3
- SPOOL:	8	3380	3
- TDISK:	12	3380	3
- USER:	80	3380	7

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES	LINESPEED
3745	3	50	56K

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- USER THINK TIME: BACTRIAN (Avg 30 Secs)
- CMS BLOCK SIZE: 4K
- USER VM SIZE: 2M
- USER RELSHARE: 100

-SERVER MACHINES:

	TYPE	NUMBER	VM SIZE	RELSHARE	OTHER OPTIONS
- VM/XA SP 2.0:	VTAM3.2/VSCS	3	64M	10000	QUICKDSP ON
	EXT. VSCS	3	64M	10000	QUICKDSP ON
- VM/ESA 1.0:	VTAM3.3/VSCS	1	64M	10000	QUICKDSP ON
	EXT. VSCS	2	64M	10000	QUICKDSP ON

NOTE: The internal VSCS has been activated but all users run through the external VSCSs.

4) MEASUREMENT SUMMARY

The following table summarizes the results of the CMS Intensive regression measurements on the 3090-600J comparing VM/XA SP 2.0 and VM/ESA 1.0. As can be seen from looking at the key indicators (response time, throughput and processor utilization), VM/ESA 1.0 has a performance improvement.

	XA SP 2.0	ESA 1.0
Storage		
REAL	512M	512M
XSTOR	2048M	2048M
Run ID	Y62R5401	Y63R5405
Variables		
USERS	5400	5400
VTAMs	3	1
VSCSs	6	2
PROCESSORS	6	6
Response Time		
TRIV INT	0.089	0.064
NONTRIV INT	0.435	0.373
TOT INT	0.278	0.227
AVG FIRST (T)	0.637	0.580
AVG LAST (T)	0.820	0.757
Throughput		
ETR (T)	187.71	189.05
ITR (H)	34.37	36.00
ITR	41.12	45.30
EMUL ITR	65.56	77.43
ITRR (H)	1.000	1.047
ITRR	1.000	1.102
Proc Usage		
PBT/CMD (H)	29.094	27.778
CP/CMD (H)	11.661	12.075
EMUL/CMD (H)	17.429	15.699
Processor Util		
TOTAL (H)	546.11	525.13
UTIL/PROC (H)	91.02	87.52
TOTEMUL (H)	327.18	296.80
TVR (H)	1.67	1.77
Storage		
WKSET (V)	53	57
PGBLPGS	96587	106K
SHRPGS	977	1079
Paging		
READS/SEC	178	265
WRITES/SEC	67	143
PAGE/CMD	1.305	2.158
XSTOR IN/SEC	1250	1286
XSTOR OUT/SEC	1396	1547
XSTOR/CMD	14.096	14.985
I/O		
VIO RATE	1962	1669
VIO/CMD	8.748	7.026
MDC READS	1133	1231
MDC WRITES	520	512
MDC MODS	411	423
MDC HIT RATIO	0.91	0.93
PRIVOPs		
PRIVOP/CMD	20.597	24.628
DIAG/CMD	17.753	17.291
DIAG 08/CMD	0.730	0.725
DIAG 10/CMD	4.816	5.787
DIAG 58/CMD	1.225	1.227
DIAG A4/CMD	4.683	4.443
DIAG A8/CMD	3.244	1.857
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 2. VM/XA SP 2.0 vs. VM/ESA 1.0 on a 3090-600J

3.2.2.3 3090-200J / 256M / VM/XA SP 2.1

Following is a description of the environment used for the CMS Intensive regression measurements on the 3090-200J comparing VM/XA SP 2.1 and VM/ESA 1.0.

```
1) WORKLOAD:     FS7B0R

2) HARDWARE CONFIGURATION

   - PROCESSOR:  3090-600J   CPU 0,1 only
   - STORAGE:
      - RSTOR:   256M
      - XSTOR:   2G

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PSYS02     3380      5
                 WKLD01     3380     16
                 WKLD02     3380     18

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - PAGE           20         3380            3
      - SPOOL           8         3380            3
      - TDISK          12         3380            3
      - USER           80         3380            7

   - COMMUNICATIONS:  CONTROLLER  NUMBER  LINES  LINESPEED
                      ----------  ------  -----  ---------
                        3745         1      50     56K

3) SOFTWARE CONFIGURATION

   - DRIVER:           TPNS
   - USER THINK TIME:  BACTRIAN (Avg 30 Secs)
   - CMS BLOCK SIZE:   4K
   - USER VM SIZE:     2M
   - USER RELSHARE:    100

   -SERVER MACHINES:    TYPE     NUMBER  VM SIZE  RELSHARE  OTHER OPTIONS
                        ----     ------  -------  --------  -------------
     - VM/XA SP 2.1: VTAM3.2/VSCS   1      64M     10000    QUICKDSP ON
     - VM/ESA 1.0:   VTAM3.3/VSCS   1      64M     10000    QUICKDSP ON

4) MEASUREMENT SUMMARY
```

The following table summarizes the results of the CMS Intensive regression measurements on the 3090-200J comparing VM/XA SP 2.1 and VM/ESA 1.0. As can be seen from looking at the key indicators (response time, throughput and processor utilization), VM/ESA 1.0 has a performance improvement.

	XA SP 2.1	ESA 1.0
Storage		
REAL	256M	256M
XSTOR	2048M	2048M
Run ID	Y2$R2001	Y23R2001
Variables		
USERS	2000	2000
VTAMs	1	1
VSCSs	2	1
PROCESSORS	2	2
Response Time		
TRIV INT	0.048	0.044
NONTRIV INT	0.353	0.296
TOT INT	0.240	0.201
AVG FIRST (T)	0.220	0.240
AVG LAST (T)	0.380	0.370
Throughput		
ETR (T)	70.38	70.90
ITR (H)	38.95	41.44
ITR	39.60	42.73
EMUL ITR	60.08	67.17
ITRR (H)	1.000	1.064
ITRR	1.000	1.079
Proc Usage		
PBT/CMD (H)	25.250	23.422
CP/CMD (H)	9.069	8.881
EMUL/CMD (H)	16.176	14.537
Processor Util		
TOTAL (H)	180.69	171.07
UTIL/PROC (H)	90.35	85.54
TOTEMUL (H)	115.75	106.19
TVR (H)	1.56	1.61
Storage		
WKSET (V)	57	82
PGBLPGS	49370	53127
SHRPGS	613	803
Paging		
READS/SEC	81	89
WRITES/SEC	1	1
PAGE/CMD	1.165	1.269
XSTOR IN/SEC	320	356
XSTOR OUT/SEC	341	395
XSTOR/CMD	9.392	10.593
I/O		
VIO RATE	743	623
VIO/CMD	10.383	8.530
MDC READS	448	480
MDC WRITES	173	173
MDC MODS	157	157
MDC HIT RATIO	0.96	0.96
PRIVOPs		
PRIVOP/CMD	25.382	20.664
DIAG/CMD	18.220	17.514
DIAG 08/CMD	0.710	0.719
DIAG 10/CMD	4.788	5.769
DIAG 58/CMD	1.222	1.213
DIAG A4/CMD	4.660	4.429
DIAG A8/CMD	3.382	1.820
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 3. VM/XA SP 2.1 vs. VM/ESA 1.0 on a 3090-200J

3.2.2.4 3090-600J / 512M / VM/XA SP 2.1

Following is a description of the environment used for the CMS Intensive regression measurements on the 3090-600J comparing VM/XA SP 2.1 and VM/ESA 1.0.

```
1) WORKLOAD:     FS7B0R

2) HARDWARE CONFIGURATION

   - PROCESSOR:  3090-600J
   - STORAGE:
      - RSTOR:   512M
      - XSTOR:   2G

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PSYS02     3380      5
                 WKLD01     3380     16
                 WKLD02     3380     18

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - PAGE           20         3380            3
      - SPOOL           8         3380            3
      - TDISK          12         3380            3
      - USER           80         3380            7

   - COMMUNICATIONS:  CONTROLLER  NUMBER  LINES  LINESPEED
                      ----------  ------  -----  ---------
                        3745         3      50     56K

3) SOFTWARE CONFIGURATION

   - DRIVER:           TPNS
   - USER THINK TIME:  BACTRIAN (Avg 30 Secs)
   - CMS BLOCK SIZE:   4K
   - USER VM SIZE:     2M
   - USER RELSHARE:    100

                       TYPE       NUMBER  VM SIZE  RELSHARE  OTHER OPTIONS
                       ----       ------  -------  --------  -------------
   - SERVER MACHINES:  VTAM/VSCS    1       64M      10000    QUICKDSP ON
                       EXT. VSCS    2       64M      10000    QUICKDSP ON
```

NOTE: The internal VSCS has been activated but all users run through the external VSCSs.

```
4) MEASUREMENT SUMMARY
```

The following table summarizes the results of the CMS Intensive regression measurements on the 3090-600J comparing VM/XA SP 2.1 and VM/ESA 1.0. As can be seen from looking at the key indicators (response time, throughput and processor utilization), VM/ESA 1.0 has a performance improvement.

	XA SP 2.1	ESA 1.0
Storage		
REAL	512M	512M
XSTOR	2048M	2048M
Run ID	Y6$R5401	Y63R5405
Variables		
USERS	5400	5400
VTAMs	3	1
VSCSs	6	2
PROCESSORS	6	6
Response Time		
TRIV INT	0.089	0.064
NONTRIV INT	0.444	0.373
TOT INT	0.283	0.227
AVG FIRST (T)	0.633	0.580
AVG LAST (T)	0.817	0.757
Throughput		
ETR (T)	188.73	189.05
ITR (H)	34.98	36.00
ITR	41.34	45.30
EMUL ITR	67.25	77.43
ITRR (H)	1.000	1.029
ITRR	1.000	1.096
Proc Usage		
PBT/CMD (H)	28.586	27.778
CP/CMD (H)	11.654	12.075
EMUL/CMD (H)	16.929	15.699
Processor Util		
TOTAL (H)	539.49	525.13
UTIL/PROC (H)	89.92	87.52
TOTEMUL (H)	319.49	296.80
TVR (H)	1.69	1.77
Storage		
WKSET (V)	51	57
PGBLPGS	96413	106K
SHRPGS	993	1079
Paging		
READS/SEC	108	265
WRITES/SEC	8	143
PAGE/CMD	0.615	2.158
XSTOR IN/SEC	1087	1286
XSTOR OUT/SEC	1156	1547
XSTOR/CMD	11.885	14.985
I/O		
VIO RATE	1975	1669
VIO/CMD	8.846	7.026
MDC READS	1151	1231
MDC WRITES	510	512
MDC MODS	420	423
MDC HIT RATIO	0.92	0.93
PRIVOPs		
PRIVOP/CMD	22.915	24.628
DIAG/CMD	17.850	17.291
DIAG 08/CMD	0.721	0.725
DIAG 10/CMD	4.827	5.787
DIAG 58/CMD	1.224	1.227
DIAG A4/CMD	4.647	4.443
DIAG A8/CMD	3.296	1.857
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 4. VM/XA SP 2.1 vs. VM/ESA 1.0 on a 3090-600J

3.3 Comparison of 370 and XA Modes

3.3.1 Introduction

A set of CMS Intensive measurements were taken to compare the performance of users running 370 virtual machines with those running XA virtual machines. The measurements were done on the 3090-600J with 5400 users giving a target processor utilization of about 90%.

3.3.2 370 to XA Measurements

Following is a description on the environment used for the CMS Intensive VM/ESA 1.0 370 Mode vs. XA Mode measurement on the 3090-600J.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

- PROCESSOR: 3090-600J
- STORAGE:
  - RSTOR: 512M
  - XSTOR: 2G

- DASD:

	PACK NAME	TYPE	CHANNEL
- SYSTEM:	PSYS02	3380	5
	WKLD01	3380	16
	WKLD02	3380	18

	NUMBER OF PACKS	TYPE	NUMBER OF CHANNELS
- PAGE	20	3380	3
- SPOOL	8	3380	3
- TDISK	12	3380	3
- USER	80	3380	7

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES	LINESPEED
3745	3	50	56K

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- USER THINK TIME: BACTRIAN (Avg 30 Secs)
- CMS BLOCK SIZE: 4K
- USER VM SIZE: 2M
- USER RELSHARE: 100

	TYPE	NUMBER	VM SIZE	RELSHARE	OTHER OPTIONS
- SERVER MACHINES:	VTAM/VSCS	1	64M	10000	QUICKDSP ON
	EXT. VSCS	2	64M	10000	QUICKDSP ON

NOTE: The internal VSCS has been activated but all users run through the external VSCSs.

4) MEASUREMENT SUMMARY

The following table gives the details of the VM/ESA 1.0 CMS Intensive measurements on the 3090-600J comparing 370 Mode and XA Mode CMS users.

A summary of the results for XA Mode appears below:

External response time took less than 0.25 seconds longer.
ITR was reduced by less than 3%.
ETR was reduced by less than 1%.
Processor utilization increased by about 2%.

	370 Mode	XA Mode
Storage		
REAL	512M	512M
XSTOR	2048M	2048M
Run ID	Y63R5405	Y63R5406
Variables		
USERS	5400	5400
VTAMs	1	1
VSCSs	2	2
PROCESSORS	6	6
Response Time		
TRIV INT	0.064	0.061
NONTRIV INT	0.373	0.346
TOT INT	0.227	0.197
AVG FIRST (T)	0.580	0.777
AVG LAST (T)	0.757	1.000
Throughput		
ETR (T)	189.05	187.79
ITR (H)	36.00	35.01
ITR	45.30	48.45
EMUL ITR	77.43	82.09
ITRR (H)	1.000	0.972
ITRR	1.000	1.070
Proc Usage		
PBT/CMD (H)	27.778	28.563
CP/CMD (H)	12.075	12.216
EMUL/CMD (H)	15.699	16.344
Processor Util		
TOTAL (H)	525.13	536.40
UTIL/PROC (H)	87.52	89.40
TOTEMUL (H)	296.80	306.93
TVR (H)	1.77	1.75
Storage		
WKSET (V)	57	57
PGBLPGS	106K	106K
SHRPGS	1079	1073
Paging		
READS/SEC	265	220
WRITES/SEC	143	132
PAGE/CMD	2.158	1.874
XSTOR IN/SEC	1286	1343
XSTOR OUT/SEC	1547	1586
XSTOR/CMD	14.985	15.597
I/O		
VIO RATE	1669	1674
VIO/CMD	7.026	6.444
MDC READS	1231	1233
MDC WRITES	512	508
MDC MODS	423	418
MDC HIT RATIO	0.93	0.93
PRIVOPs		
PRIVOP/CMD	24.628	24.674
DIAG/CMD	17.291	17.599
DIAG 08/CMD	0.725	0.735
DIAG 10/CMD	5.787	5.996
DIAG 58/CMD	1.227	1.225
DIAG A4/CMD	4.443	4.473
DIAG A8/CMD	1.857	1.901
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 5. CMS Intensive - 370 Mode vs. XA Mode

3.4 EDF and SFS Comparisons

3.4.1 Introduction

The measurements in this section compare the performance of the Shared File System (SFS) to the CMS minidisk file system (EDF). All measurements were done on ESA 1.0. Two comparisons are made: one at approximately equal processor utilization and the other with an equal number of users.

There were four production SFS file pools and two read only SFS file pools. Almost all of the SFS activity was with the production file pools. Four production file pools were chosen in order to minimize the likelihood of rollbacks due to deadlock. 'Rollbacks due to deadlock' was the limiting factor for these measurements that made it impractical to have fewer than four file pools. Other possible factors that can limit the number of users supportable by a single file pool include:

- checkpoint serialization
- file pool recovery time
- uniprocessor server (The SFS server can run on only one real processor at a time.)

The two read only file pools contained read only files which were shared across users. An installation may wish to put such files in a single, read only file pool. There were two here for added flexibility for doing performance measurements.

SFS counts and timings obtained from the QUERY FILEPOOL STATUS command are provided in "Appendix B. SFS Counter Data" on page 155 to supplement the SFS measurement data provided in this section.

3.4.2 Measurements

3.4.2.1 3090-600J / 512M / EDF and SFS at Equal Utilization

This section summarizes the results of measurements that compare ESA 1.0 EDF and ESA 1.0 SFS at the same processor utilization.

1) WORKLOAD: FS7B0R and FS7B35R

2) HARDWARE CONFIGURATION

- PROCESSOR: 3090-600J
- STORAGE:
  - RSTOR: 512M
  - XSTOR: 2G

- DASD:

	PACK NAME	TYPE	CHANNEL
- SYSTEM:	PSPT01	3380	5
	WKLD01	3380	16
	WKLD02	3380	18

	NUMBER OF PACKS	TYPE	NUMBER/TYPE OF CHANNELS
- PAGE:	20	3380-AA4	3 - 3880-3
- SPOOL:	8	3380-AA4	3 - 3880-3
- TDISK:	12	3380-AA4	3 - 3880-3
- USER:	18	3380-AK4	4 - 3990-2

- TAPE: MONITOR 3480 CHANNEL C

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES	LINESPEED
3745	3	50	56K

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- USER THINK TIME: BACTRIAN (Avg 30 Secs)
- CMS BLOCK SIZE: 4K
- USER VM SIZE: 2M
- USER RELSHARE: 100

	TYPE	NUMBER	VM SIZE	RELSHARE	OTHER OPTIONS
- SERVER MACHINES:	VTAM	1	64M	10000	QUICKDSP ON
	EXT. VSCS	2	64M	10000	QUICKDSP ON
	SFS (prod)	4	32M	1500	QUICKDSP ON
	SFS (R/O)	2	32M	1500	QUICKDSP ON
	CRR	1	16M	1500	QUICKDSP ON

4) MEASUREMENT SUMMARY

The target processor utilization for both the 0% SFS (EDF) run and the 35% SFS run was 90%. EDF supported 5700 users and SFS supported 4800 users (a 16 percent reduction).

The processor busy time per command (PBT/CMD (H)) increased by 14.2%. The increase for the corresponding ESA 1.0 370 Feature measurements averaged 17.6% (see “ 4381-13 / 16M / EDF and SFS at Equal Utilization” on page 101). The ESA 1.0 Feature showed a smaller percentage increase than ESA 1.0 370 Feature because:

- VTAM processing for the ESA 1.0 Feature measurements increases the base. To illustrate this, if we subtract the VTAM processor busy time per command (TOT CPU/CMD (V) under VTAM Machines) from the ESA 1.0 Feature times, the ESA 1.0 Feature percentage increase becomes 17.1%.

- Minidisk cache works more efficiently with SFS than with EDF. See " 3090-600J / 512M / EDF and SFS with Equal Users" on page 55.
- The higher SFS real storage requirements had less impact because most of the paging was to expanded storage.

For the ESA 1.0 Feature, the CP processor time per command (CP/CMD (H)) accounted for 55% of the increased processor time per command when comparing EDF with SFS. In ESA 1.0 370 Feature, the CP processor time per command only accounted for 39% of the increased processor time. This shift of the SFS CPU delta toward CP is because of higher APPC and block IO pathlengths with the ESA 1.0 Feature. This would have tended to make the gap between EDF and SFS larger on the ESA 1.0 Feature, but it was outweighed by the three items listed above.

Virtual file I/Os per command increased by 7% (from 5.86 for EDF to 6.28 for SFS). The virtual file I/Os per command can be approximated by adding:

- DIAG A4/CMD
- DIAG A8/CMD (this included some additional I/O that is not file system related, but it should be about the same in both cases)
- IO/CMD (Q) (for the 35% SFS run)

The sum of PAGE/CMD and XSTOR/CMD is a good measure of contention for real storage. For EDF the sum is 19.4 and for SFS the sum is 21.5. This increase is similar to the ESA 1.0 370 Feature paging increase shown in " 4381-13 / 16M / EDF and SFS at Equal Utilization" on page 101. In a comparison at equal processor utilization, overall real storage contention is only slightly higher in the SFS case because most of SFS's increased real storage requirement is offset by the storage that is made available by running with fewer users (4800 instead of 5700).

These equal utilization measurements show that the total internal response times (TOT INT) are similar for EDF and SFS. The RTM command rates were not consistent, therefore response times were renormalized over the number of commands recorded by TPNS. After renormalizing, the response time for EDF is 0.436 seconds and for SFS it is 0.413 seconds.

The external response times (AVG LAST (T)) are 1.560 seconds for EDF and only 0.977 seconds for SFS. Since the internal response times are similar, the higher EDF external response time is due to the greater VTAM and network contention caused by the 5700 users (the SFS run only had 4800 users).

	VM/ESA 1.0 EDF	VM/ESA 1.0 35% SFS
CPU Util. (approx)	90	90
Storage		
REAL	512M	512M
XSTOR	2048M	2048M
Run ID	Y63R5705	Y63F4804
Variables		
USERS	5700	4800
VTAMs	1	1
VSCSs	2	2
PROCESSORS	6	6
Response Time		
TRIV INT	0.071	0.063
NONTRIV INT	0.550	0.557
TOT INT	0.294	0.326
AVG FIRST (T)	1.183	0.725
AVG LAST (T)	1.560	1.000
Throughput		
ETR (T)	194.13	166.99
ITR (H)	35.58	31.12
ITR	52.85	39.53
EMUL ITR	91.18	70.43
ITRR (H)	1.000	0.875
ITRR	1.000	0.748
Proc Usage		
PBT/CMD (H)	28.11	32.14
CP/CMD (H)	12.36	14.62
EMUL/CMD (H)	15.75	17.52
Processor Util		
TOTAL (H)	545.69	536.68
UTIL/PROC (H)	90.95	89.45
TOTEMUL (H)	305.70	292.49
TVR (H)	1.78	1.83
Storage		
WKSET (V)	51	56
PGBL PAGES	104K	108K
SHAR PAGES	1220	1353
Paging		
READS/SEC	340	223
WRITES/SEC	237	145
PAGE/CMD	2.972	2.204
XSTOR IN/SEC	1431	1509
XSTOR OUT/SEC	1761	1718
XSTOR/CMD	16.442	19.324
I/O		
VIO RATE	1628	1127
VIO/CMD	5.653	5.324
MDC READS	1134	989
MDC WRITES	511	305
MDC MODS	434	238
MDC HIT RATIO	0.93	0.93
PRIVOPs		
PRIVOP/CMD	29.396	41.149
DIAG/CMD	20.335	18.202
DIAG 08/CMD	0.721	0.737
DIAG 10/CMD	5.192	4.803
DIAG 58/CMD	1.221	1.234
DIAG A4/CMD	3.966	2.599
DIAG A8/CMD	1.890	1.623

	VM/ESA 1.0 EDF	VM/ESA 1.0 35% SFS
CPU Util. (approx)	90	90
Storage		
REAL	512M	512M
XSTOR	2048M	2048M
Run ID	Y63R5705	Y63F4804
Variables		
USERS	5700	4800
VTAMs	1	1
VSCSs	2	2
PROCESSORS	6	6
VTAM Machines		
WKSET (V)	3348	2841
TOT CPU/CMD (V)	0.0053	0.0054
CP CPU/CMD (V)	0.0034	0.0035
EMUL/CMD (V)	0.0019	0.0019
SFS Servers		
WKSET (V)	na	3236
TOT CPU/CMD (V)	na	0.0045
CP CPU/CMD (V)	na	0.0024
EMUL/CMD (V)	na	0.0020
FP REQ/CMD(Q)	na	1.34
IO/CMD (Q)	na	2.06
IO TIME/CMD (Q)	na	0.05
SFS TIME/CMD (Q)	na	0.15
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 6. 3090-600J / 512M / EDF and SFS at Equal Utilization

3.4.2.2 3090-600J / 512M / EDF and SFS with Equal Users

This section summarizes the results of measurements that compare ESA 1.0 EDF and ESA 1.0 SFS at the same number of users.

1) WORKLOAD: FS7B0R and FS7B35R

2) HARDWARE CONFIGURATION

- PROCESSOR: 3090-600J
- STORAGE:
  - RSTOR: 512M
  - XSTOR: 2G

- DASD:

	PACK NAME	TYPE	CHANNEL
- SYSTEM:	PSPT01	3380	5
	WKLD01	3380	16
	WKLD02	3380	18

	NUMBER OF PACKS	TYPE	NUMBER/TYPE OF CHANNELS
- PAGE:	20	3380-AA4	3 - 3880-3
- SPOOL:	8	3380-AA4	3 - 3880-3
- TDISK:	12	3380-AA4	3 - 3880-3
- USER:	18	3380-AK4	4 - 3990-2

- TAPE: MONITOR 3480 CHANNEL C

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES	LINESPEED
3745	3	50	56K

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- USER THINK TIME: BACTRIAN (Avg 30 Secs)
- CMS BLOCK SIZE: 4K
- USER VM SIZE: 2M
- USER RELSHARE: 100

	TYPE	NUMBER	VM SIZE	RELSHARE	OTHER OPTIONS
- SERVER MACHINES:	VTAM	1	64M	10000	QUICKDSP ON
	EXT. VSCS	2	64M	10000	QUICKDSP ON
	SFS (prod)	4	32M	1500	QUICKDSP ON
	SFS (R/O)	2	32M	1500	QUICKDSP ON
	CRR	1	16M	1500	QUICKDSP ON

4) MEASUREMENT SUMMARY

Storage usage, processor usage, and response times all showed increases compared to EDF when holding the number of users constant. The 0% SFS (EDF) run and the 35% SFS run both had 4800 users.

The sum of PAGE/CMD and XSTOR/CMD is a good measure of contention for real storage. For EDF the sum is 13.1 and for SFS the sum is 21.5. This increase is similar to the ESA 1.0 370 Feature storage increase shown in “ 4381-13 / 16M / EDF and SFS with Equal Users” on page 104, and is indicative of SFS’s higher real storage requirements.

The processor busy time per command (PBT/CMD (H)) showed an increase of 17%. This increase is greater than the 14.2% increase observed for the equal utilization comparison, primarily because of the higher paging.

Virtual file I/Os per command increased by 9.4% (from 5.74 for EDF to 6.28 for SFS).

For both the equal utilization and equal users measurements for the ESA 1.0 Feature, the minidisk cache (MDC) is equally effective for EDF and SFS at reducing DASD read I/Os, as evidenced by similar MDC HIT RATIOs. The significantly reduced MDC MODS and MDC WRITES rates show that CP/ESA manages the cache more efficiently in the SFS case. This is because with SFS, much of the write activity that is in support of file directory updates is to the SFS logs, which (since they are nearly write-only) are made ineligible for MDC. This optimization is not feasible in the EDF case because each minidisk has its own directory contained within it. There are two additional reasons why SFS and MDC work well together:

1. SFS uses block IO, which has a special synchronous path when all the data requested is available in the minidisk cache. This avoids the extra pathlength associated with asynchronous processing. EDF doesn't use block IO, so this improvement does not apply.

2. SFS buffers are always 4K aligned-- the case that MDC handles the most efficiently. With EDF, the user is not required to have a 4K aligned buffer.

These equal user measurements show that the total internal response time (TOT INT) is longer for SFS. The RTM command rates were not consistent, therefore response times were renormalized over the number of commands recorded by TPNS. After renormalizing, the response time for EDF is 0.218 seconds and for SFS it is 0.413 seconds. This increase is due to SFS running at a higher level of processor utilization and at a higher level of real storage contention. As shown in the previous section, when SFS is run at a similar level of contention as EDF, response times are comparable.

	VM/ESA 1.0 EDF	VM/ESA 1.0 35% SFS
Storage		
REAL	512M	512M
XSTOR	2048M	2048M
Run ID	Y63R4801	Y63F4804
Variables		
USERS	4800	4800
VTAMs	1	1
VSCSs	2	2
PROCESSORS	6	6
Response Time		
TRIV INT	0.052	0.063
NONTRIV INT	0.300	0.557
TOT INT	0.199	0.326
AVG FIRST (T)	0.325	0.725
AVG LAST (T)	0.455	1.000
Throughput		
ETR (T)	169.75	166.99
ITR (H)	36.43	31.12
ITR	39.98	39.53
EMUL ITR	69.19	70.43
ITRR (H)	1.000	0.854
ITRR	1.000	0.989
Proc Usage		
PBT/CMD (H)	27.45	32.14
CP/CMD (H)	12.05	14.62
EMUL/CMD (H)	15.40	17.52
Processor Util		
TOTAL (H)	465.98	536.68
UTIL/PROC (H)	77.66	89.45
TOTEMUL (H)	261.42	292.49
TVR (H)	1.78	1.83
Storage		
WKSET (V)	58	56
PGBL PAGES	108K	108K
SHAR PAGES	1115	1353
Paging		
READS/SEC	220	223
WRITES/SEC	109	145
PAGE/CMD	1.938	2.204
XSTOR IN/SEC	873	1509
XSTOR OUT/SEC	1029	1718
XSTOR/CMD	11.205	19.324
I/O		
VIO RATE	1401	1127
VIO/CMD	7.528	5.324
MDC READS	965	989
MDC WRITES	435	305
MDC MODS	374	238
MDC HIT RATIO	0.94	0.93
PRIVOPs		
PRIVOP/CMD	33.149	41.149
DIAG/CMD	20.847	18.202
DIAG 08/CMD	0.719	0.737
DIAG 10/CMD	5.108	4.803
DIAG 58/CMD	1.231	1.234
DIAG A4/CMD	3.853	2.599
DIAG A8/CMD	1.885	1.623

	VM/ESA 1.0 EDF	VM/ESA 1.0 35% SFS
Storage		
REAL	512M	512M
XSTOR	2048M	2048M
Run ID	Y63R4801	Y63F4804
Variables		
USERS	4800	4800
VTAMs	1	1
VSCSs	2	2
PROCESSORS	6	6
VTAM Machines		
WKSET (V)	2762	2841
TOT CPU/CMD (V)	0.0056	0.0054
CP CPU/CMD (V)	0.0036	0.0035
EMUL/CMD (V)	0.0020	0.0019
SFS Servers		
WKSET (V)	na	3236
TOT CPU/CMD (V)	na	0.0045
CP CPU/CMD (V)	na	0.0024
EMUL/CMD (V)	na	0.0020
FP REQ/CMD(Q)	na	1.34
IO/CMD (Q)	na	2.06
IO TIME/CMD (Q)	na	0.05
SFS TIME/CMD (Q)	na	0.15
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 7. 3090-600J/ 512M / EDF and SFS with Equal Users

3.5 Connectivity

3.5.1 Introduction

This section describes the local APPC/VM and AVS measurements made for ESA 1.0.

VMPRF, RTM/SF, FSID, and the Connectivity Driver tool were used to produce this data. In all the tables in this section, items of the form ″xxx/CMD″ (for example, ″PBT/CMD″) have a slightly different meaning than similar variables elsewhere. In all cases, ″CMD″ here means not one user interaction, but one APPCVM SEND/RECEIVE pair. A single user interaction will, in general, cause several of these SEND/RECEIVE pairs to occur.

3.5.2 Local APPC Measurements

The performance of local APPC on ESA 1.0 was compared to local APPC on HPO 5.0 in two environments: *single-user* and *multi-user*. The single-user measurement consists of running the connectivity driver tool (see “Connectivity Driver Tool” on page 176 for a description of this tool) in two virtual machines. One machine is the *user* and the other is the *resource*. A large number of APPC/VM transfers is made and the processor busy time is noted. This benchmark gives a fairly good indication of the pure APPC/VM performance of the system.

The multi-user environment involves running many user machines who communicate with a small number of resource machines. Each user repeatedly issues a few APPC/VM transfers followed by a short wait. This benchmark gives an indication of the APPC/VM throughput of the system.

3.5.2.1 APPC/VM Single-user

```
1) WORKLOAD:     CONN1

2) HARDWARE CONFIGURATION

   - PROCESSOR:  3090-300J CPU 0 only, or CPU 0,1 only
   - STORAGE:
      - RSTOR:   64M
      - XSTOR:   512M

   - DASD:       No DASD I/O in this benchmark

   - TAPE:       MONWRITE 3480
```

```
3) SOFTWARE CONFIGURATION

   - DRIVER:           Manual
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)
   - USER RELSHARE:    100 (Users on ESA 1.0 )
   - USER RELSHARE:    10000 (Resources on ESA 1.0 )


   - HPO 5.0
                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
     - SERVER MACHINE: Resource  1 per      3M       1     QDROP OFF USERS
                                 processor                   FAVORED
                                                             FAVORED 100

   - ESA 1.0
                       TYPE      NUMBER  VM SIZE  RELSHARE  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
     - SERVER MACHINE: Resource  1 per      3M     10000    QUICKDSP ON
                                 processor

4) MEASUREMENT SUMMARY
```

Table 8 on page 61 summarizes the results of some single-user local APPC runs. For each measurement, an APPC/VM connection was established between a user and a resource virtual machine. The user then issued a large number of transactions without pausing; in the 512- and 4096-byte cases this number was 300,000 and in the 102400-byte case it was 100,000. Figure 9 on page 15 and Figure 10 on page 16 present the data in this table graphically. The benchmark was run twice on HPO 5.0 and twice on ESA 1.0 to examine the effect of running with more than one real processor. In all cases there is only one resource virtual machine and one user virtual machine. Both HPO 5.0 and ESA 1.0 use more processor time in an MP environment (except ESA 1.0 in the 102400-byte case), but ESA 1.0 degrades less than HPO 5.0.

The degradation observed going from HPO 5.0 to ESA 1.0 is mostly due to increased CP pathlength in ESA 1.0. For this benchmark, the increased pathlength is mostly IUCV and module linkage code. In the following tables, *application instrumentation* refers to counters in the Connectivity Driver application that record IUCV transactions. All times are in milliseconds.

RELEASE	HPO 5.0		ESA 1.0	
Storage REAL STORAGE EXPANDED STOR	64M 512M		64M 512M	
Run-ID	G0350	G0350	G1622	G1622
Variables USERS PROCESSORS	1 1	1 2	1 1	1 2
512 byte PBT/CMD VIRT/CMD CP/CMD	 0.410 0.035 0.375	 0.568 0.036 0.532	 0.616 0.053 0.559	 0.690 0.049 0.641
4096 byte PBT/CMD VIRT/CMD CP/CMD	 0.433 0.037 0.397	 0.594 0.035 0.558	 0.644 0.051 0.593	 0.733 0.049 0.684
102400 byte PBT/CMD VIRT/CMD CP/CMD	 1.541 0.059 1.483	 1.716 0.057 1.659	 2.564 0.074 2.492	 2.445 0.074 2.371
Note: All data derived from VMPRF and application instrumentation.				

Table 8. Local APPC/VM single-user: HPO 5.0 vs. ESA 1.0 on one and two 3090-300J processors

3.5.2.2 APPC/VM Multi-user

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   - PROCESSOR:  3090-300J CPU 0,1 only
   - STORAGE:
      - RSTOR:   16M (ESA 1.0 )
                 64M (HPO 5.0 )
      - XSTOR:   512M

   - DASD:       No DASD I/O in this benchmark

   - TAPE:       MONWRITE 3480

3) SOFTWARE CONFIGURATION

   - DRIVER:          FSID
   - USER THINK TIME: RANDOM (Avg 1 Secs)
   - REQUEST SIZE 1:  512 bytes
   - TRANSACTION 1:   11 requests of 512 bytes each
   - REQUEST SIZE 2:  102400 bytes
   - TRANSACTION 2:   22 requests of 102400 bytes each
   - CMS BLOCKSIZE:   4K
   - USER VM SIZE:    2M
   - USER PRIORITY:   64 (Users on HPO 5.0 )
   - USER PRIORITY:   1 (Resources on HPO 5.0 )
   - USER RELSHARE:   100 (Users on ESA 1.0 )
   - USER RELSHARE:   10000 (Resources on ESA 1.0 )


   - HPO 5.0
                        TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                        ----      ------  -------  --------  -------------
     - SERVER MACHINE: Resource    2        3M       1      QDROP OFF USERS
                                                            FAVORED
                                                            FAVORED 100

   - ESA 1.0
                        TYPE      NUMBER  VM SIZE  RELSHARE  OTHER OPTIONS
                        ----      ------  -------  --------  -------------
     - SERVER MACHINE: Resource    2        3M     10000    QUICKDSP ON


4) MEASUREMENT SUMMARY
```

Table 9 on page 63 summarizes the results of two multi-user local APPC runs done to compare the throughput of APPC/VM on ESA 1.0 to the throughput on HPO 5.0. Each measurement was made on 2 processors of a 3090-300J; the third processor was physically offline. All these runs involve two resource virtual machines. The 512-byte runs measure 150 user machines and the 102400-byte runs have 30 users. After all the user machines logon and connect to one or the other resource machine, they repeatedly issue a small number of APPC/VM transactions and then pause for 1 second. Figure 11 on page 17presents the data in this table graphically.

As in the single-user benchmark, the degradation observed going from HPO 5.0 to ESA 1.0 is mostly due to increased CP pathlength in ESA 1.0. Again, the increased pathlength is mostly IUCV and module linkage code.

In all this data, any quantity given as ″per command″ means ″per single APPC request.″ For example, a transaction consisting of 11 512-byte requests counts as 11 ″commands.″ All times are in milliseconds.

	HPO 5.0		ESA 1.0	
Storage REAL STORAGE EXPANDED STOR	64M 512M		16M 512M	
Run-ID	G1552	G1552	G1730	G1730
Variables USERS PROCESSORS TRANSACTION	150 2 512	30 2 102400	150 2 512	30 2 102400
Transaction	512	102400	512	102400
Throughput CMDS/SEC	1243.7	442.0	1343.7	451.9
CPU Usage PBT/CMD VIRT/CMD CP/CMD	0.81 0.05 0.75	2.28 0.07 2.21	0.95 0.07 0.87	2.95 0.09 2.85
CPU Utilization TOTCPU CPCPU VIRTCPU	100.4 94.1 6.4	100.7 97.7 3.0	127.3 117.2 10.1	133.1 129.1 4.0
Paging PAGE RATE PAGE/CMD	0.0 0.0	0.60 0.0014	0.0 0.0	0.00 0.000
I/O VIO RATE VIO/CMD	226.2 0.182	42.4 0.0955	244.0 0.182	41.0 0.091
Note: All HPO 5.0 data is from VMMAP and application instrumentation. ESA 1.0 processor utilization data is from VMPRF and application instrumentation. ESA 1.0 paging and VIO data is from RTM/SF.				

Table 9. Local APPC/VM Multi-user: 3090-300J (two processors only)

3.5.3 AVS Measurements

This section describes some multi-user AVS runs made running on ESA 1.0 These runs are not compared with any HPO 6.0 environments; they were made only to study the behavior of AVS under various loads.

3.5.3.1 AVS Multi-user

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   USER SIDE
   - PROCESSOR:  3090-600J
   - STORAGE:
      - RSTOR:   512M
      - XSTOR:   2G

   - DASD:       No DASD I/O in this benchmark

   - TAPE:       MONWRITE 3480

   RESOURCE SIDE
   - PROCESSOR:  3090-300S
   - STORAGE:
      - RSTOR:   256M
      - XSTOR:   512M

   - DASD:       No DASD I/O in this benchmark

   - TAPE:       MONWRITE 3480

3) SOFTWARE CONFIGURATION

   - DRIVER:          FSID
   - USER THINK TIME: RANDOM (Avg 10 Secs)
   - REQUEST SIZE 1:  512 bytes
   - TRANSACTION 1:   11 requests of 512 bytes each
   - REQUEST SIZE 2:  102400 bytes
   - TRANSACTION 2:   22 requests of 102400 bytes each
   - CMS BLOCKSIZE:   4K
   - USER VM SIZE:    2M
   - USER RELSHARE:   100
   - USER RELSHARE:   10000


                       TYPE      NUMBER  VM SIZE  RELSHARE  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   VTAM/AVS     3      64M     10000    QUICKDSP ON
                       Resource  1 per      3M     10000    QUICKDSP ON
                                 processor

4) MEASUREMENT SUMMARY
```

Table 10 on page 65 summarizes the results of two multi-user AVS runs done to study the throughput and cost per transaction of AVS/XA on ESA 1.0. AVS/XA and VTAM 3.3 were run together in a 64M XA-mode virtual machine. The think time, transaction size and number of active users were chosen to achieve a reasonable level of processor and 3088 usage. The first two data columns on this chart refer to a run in which the data transfer size was 512 bytes. The last two columns describe a 102400-byte run.

In all this data, any quantity given as ″per command″ means ″per single APPC request.″ For example, a transaction consisting of 11 512-byte requests counts as 11 ″commands.″

	ESA 1.0			
Storage				
REAL STORAGE	512M	256M	512M	256M
EXPANDED STOR	2G	512M	2G	512M
Run-ID	1		2	
Variables				
USERS	2000	3	120	6
PROCESSORS	6	3	6	3
AVS/VTAM MACH.	3	3	3	3
TRANSACTION	512	512	102400	512
SIDE	User	Resource	User	Resource
Throughput				
CMDS/SEC	606.0	na	30.8	na
CPU Usage				
PBT/CMD	0.00340	0.00370	0.05900	0.02820
VIRT/CMD	0.00166	0.00144	0.01135	0.00810
CP/CMD	0.00174	0.00226	0.04765	0.02010
CPU Utilization				
TOTCPU	206	224	182	87
CPCPU	101	87	35	25
VIRTCPU	105	137	147	62
Paging				
PAGE RATE	0.0	0.0	0.0	0.0
PAGE/CMD	0.0	0.0	0.0	0.0
I/O				
VIO RATE	104	0	2	0
VIO/CMD	0.172	0	0.064	0
VTAM/AVS				
PBT/CMD	0.0026	0.0029	0.0564	0.0268
TOTCPU	159.7	188.2	173.7	82.7
Note: All data is from RTM/SF and application instrumentation.				

Table 10. AVS Multi-user runs: 3090-600J to 3090-300S via 3088

3.6 Processor Capacity

3.6.1 Introduction

The processor capacity study is a set of measurements of VM/ESA 1.0 on different 3090J processors. The study was on the 180J, 200J, 300J and 600J. The way the measurements were obtained was to take a 3090-600J and to vary off the number of processors needed to have a 180, 200, 300 and 600 configuration. As the number of processors was increased, the number of users was increased in order to keep the average processor utilization at approximately 90%. Measurements were made with the CMS intensive workload (FS7B0R).

3.6.2 CMS Intensive Measurements

3.6.2.1 VM/ESA 1.0 / 90% Processor Utilization

Following is a description of the environment used for the CMS Intensive processor capacity measurements.

```
1) WORKLOAD:     FS7B0R

2) HARDWARE CONFIGURATION

   - PROCESSOR:
      - 180J:    3090-600J    CPU 1 only
      - 200J:    3090-600J    CPU 0,1 only
      - 300J:    3090-600J    CPU 0,1,2 only
      - 600J:    3090-600J

   - STORAGE:
      - 180J:
        - RSTOR:   128M
        - XSTOR:   256M
      - 200J:
        - RSTOR:   256M
        - XSTOR:     2G
      - 300J:
        - RSTOR:   256M
        - XSTOR:     2G
      - 600J:
        - RSTOR:   512M
        - XSTOR:     2G

   - DASD:        PACK NAME  TYPE  CHANNEL
                  ---------  ----  -------
      - SYSTEM:   PSYS02     3380      5
                  WKLD01     3380     16
                  WKLD02     3380     18

                  NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                  ---------------  ----  ------------------
      - PAGE:           20         3380            3
      - SPOOL:           8         3380            3
      - TDISK:          12         3380            3
      - USER:           80         3380            7

   - COMMUNICATIONS:  CONTROLLER  NUMBER  LINES  LINESPEED
                      ----------  ------  -----  ---------
     - 180J             3745        1      50     56K
     - 200J             3745        1      50     56K
     - 300J             3745        3      50     56K
     - 600J             3745        3      50     56K
```

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- USER THINK TIME: BACTRIAN (Avg 30 Secs)
- CMS BLOCK SIZE: 4K
- USER VM SIZE: 2M
- USER RELSHARE: 100

- SERVER MACHINES:	TYPE	NUMBER	VM SIZE	RELSHARE	OTHER OPTIONS
- 180J:	VTAM/VSCS	1	64M	10000	QUICKDSP ON
- 200J:	VTAM/VSCS	1	64M	10000	QUICKDSP ON
- 300J:	VTAM/VSCS	1	64M	10000	QUICKDSP ON
	EXT. VSCS	2	64M	10000	QUICKDSP ON
- 600J:	VTAM/VSCS	1	64M	10000	QUICKDSP ON
	EXT. VSCS	2	64M	10000	QUICKDSP ON

NOTE: On the 300J and the 600J, the internal VSCS has been activated but all users run through the external VSCSs.

4) MEASUREMENT SUMMARY

The following table details the results of the CMS Intensive Processor study. Each column shows the growth in the number of users and the other data gathered for these measurements.

> Examination of the data shows that performance of the measurement on the 3090-600J is out of
> line with the data for the smaller processors. This artificial degradation is the result of the use of
> the SET AUTOREAD ON command in the test scripts for command synchronization by the
> workload driver (TPNS). This causes users to enter a VMREAD state after completing commands.
> When the user enters VMREAD state its virtual machine issues an IUCV two-way send to the
> VSCS virtual machine and is placed on the RECEIVE queue which is a FIFO queue. The user's
> message block stays on this queue until a reply comes, causing the queue to be searched to match
> the reply with the appropriate user's message block. Since the queue is FIFO users with longer
> think times migrate to the top of the queue, causing a search of the entire queue in most instances.
> The increase in the number of users highlights this problem on the 6-way.

> To improve performance for the case when users are in a VMREAD state HCPIUF was changed
> so that the RECEIVE queue is now LIFO instead of FIFO. This causes users with long think
> times to migrate toward the bottom while those whose requests complete in a shorter timeframe
> can be found near the top. Now the entire queue is rarely searched giving an overall improvement
> in performance. As a result performance is more in line with expectations as evidenced by the data
> below comparing ESA 1.0 at the starter set level to the same level plus the aforementioned fix to
> HCPIUF. The data presented for ESA 1.0 with the HCPIUF modifications should be much closer
> to actual customer expectations for a 3090-600J than the ESA 1.0 data without the changes to
> HCPIUF.

> Please note that the data below for the modified HCPIUF code is for prototype code. Although
> these measurements have nearly 100% of the users on VSCS's IUCV RECEIVE queue, sampled
> installations had only about 10% of their users in this queue. Since it was determined that this
> modification would not greatly benefit real users the change won't be available in ESA 1.0 but has
> been rolled into ESA 1.1. Only those SNA installations whose users enter VMREAD state fre-
> quently will see an improvement by applying this fix. Installations that have only a small fraction
> of their users in VMREAD state will see no change in their performance.

	180J	200J	300J	600J
Storage				
REAL	128M	256M	256M	512M
XSTOR	256M	2048M	2048M	2048M
Run ID	Y13R1001	Y23R2101	Y33R3001	Y63R5601
Variables				
USERS	1000	2100	3000	5600
VTAMs	1	1	1	1
VSCSs	1	1	2	2
PROCESSORS	1	2	3	6
Response Time				
TRIV INT	0.045	0.044	0.053	0.066
NONTRIV INT	0.392	0.325	0.326	0.429
TOT INT	0.267	0.219	0.219	0.245
AVG FIRST (T)	0.310	0.250	0.295	0.807
AVG LAST (T)	0.500	0.400	0.435	1.033
Throughput				
ETR (T)	35.88	74.39	107.25	193.95
ITR (H)	41.15	41.33	39.55	36.34
ITR	43.77	42.88	42.27	49.64
EMUL ITR	69.63	67.61	68.64	86.18
ITRR (H)	1.000	1.004	0.961	0.883
ITRR	1.000	0.980	0.966	1.134
Proc Usage				
PBT/CMD (H)	22.858	23.368	23.669	27.521
CP/CMD (H)	8.753	8.902	9.515	12.198
EMUL/CMD (H)	14.101	14.463	14.149	15.319
Processor Util				
TOTAL (H)	87.18	180.00	271.19	533.77
UTIL/PROC (H)	87.18	90.00	90.40	88.96
TOTEMUL (H)	53.78	111.40	162.12	297.09
TVR (H)	1.62	1.62	1.67	1.80
Storage				
WKSET (V)	74	72	62	53
PGBLPGS	25432	52640	49540	105K
SHRPGS	683	794	888	1156
Paging				
READS/SEC	212	84	115	242
WRITES/SEC	144	1	26	140
PAGE/CMD	9.923	1.143	1.315	1.970
XSTOR IN/SEC	82	388	600	1408
XSTOR OUT/SEC	250	429	675	1611
XSTOR/CMD	9.254	10.983	11.888	15.566
I/O				
VIO RATE	322	655	952	1624
VIO/CMD	8.443	8.503	8.309	6.152
MDC READS	229	503	709	1114
MDC WRITES	95	183	274	519
MDC MODS	78	165	240	428
MDC HIT RATIO	0.93	0.96	0.95	0.92
PRIVOPs				
PRIVOP/CMD	20.937	20.392	26.595	24.581
DIAG/CMD	18.083	17.191	17.607	16.525
DIAG 08/CMD	0.753	0.712	0.727	0.722
DIAG 10/CMD	5.798	5.767	5.772	5.171
DIAG 58/CMD	1.226	1.223	1.231	1.217
DIAG A4/CMD	4.376	4.436	4.410	3.944
DIAG A8/CMD	2.007	1.842	1.921	1.908
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Table 11. 3090J CMS Intensive Processor Capacity Measurements

	ESA 1.0	ESA 1.0	DELTA	% DIFF
Storage				
REAL	512M	512M	0M	0.00%
XSTOR	2048M	2048M	0M	0.00%
Run ID	Y63R5709	Y63R570C	DELTA	% DIFF
Variables				
USERS	5700	5700	0	0.00%
VTAMs	1	1	0	0.00%
VSCSs	2	2	0	0.00%
PROCESSORS	6	6	0	0.00%
Response Time				
TRIV INT	0.066	0.065	-0.001	-1.52%
NONTRIV INT	0.476	0.344	-0.132	-27.73%
TOT INT	0.259	0.222	-0.037	-14.29%
AVG FIRST (T)	1.200	0.459	-0.741	-61.75%
AVG LAST (T)	1.547	0.609	-0.938	-60.63%
Throughput				
ETR (T)	193.90	199.51	5.61	2.89%
ITR (H)	214.97	231.75	16.78	7.81%
ITR	51.48	44.55	-6.93	-13.46%
EMUL ITR	88.97	72.71	-16.26	-18.28%
ITRR (H)	1.000	1.078	0.078	7.80%
ITRR	1.000	0.865	-0.135	-13.50%
Proc Usage				
PBT/CMD (H)	27.911	25.890	-2.021	-7.24%
PBT/CMD	27.901	25.863	-2.038	-7.30%
CP/CMD (H)	12.261	10.479	-1.782	-14.53%
CP/CMD	11.759	10.025	-1.734	-14.75%
EMUL/CMD (H)	15.646	15.408	-0.238	-1.52%
EMUL/CMD	16.143	15.839	-0.304	-1.88%
Processor Util				
TOTAL (H)	541.18	516.54	-24.64	-4.55%
TOTAL	541	516	-25	-4.62%
UTIL/PROC (H)	90.20	86.09	-4.11	-4.56%
UTIL/PROC	90.17	86.00	-4.17	-4.62%
TOTEMUL (H)	303.38	307.39	4.01	1.32%
TOTEMUL	313	316	3	0.96%
TVR (H)	1.78	1.68	-0.10	-5.62%
TVR	1.73	1.63	-0.10	-5.78%
Storage				
WKSET (V)	88	55	-33	-37.50%
PGBLPGS	104K	104K	0K	0.00%
SHRPGS	1158	1203	45	3.89%
Paging				
READS/SEC	250	245	-5	-2.00%
WRITES/SEC	147	126	-21	-14.29%
PAGE/CMD	2.047	1.860	-0.187	-9.14%
XSTOR IN/SEC	1127	1011	-116	-10.29%
XSTOR OUT/SEC	1322	1202	-120	-9.08%
XSTOR/CMD	12.630	11.092	-1.538	-12.18%
I/O				
VIO RATE	1630	1666	36	2.21%
VIO/CMD	8.407	8.350	-0.057	-0.68%
MDC READS	1131	1148	17	1.50%
MDC WRITES	522	529	7	1.34%
MDC MODS	437	445	8	1.83%
MDC HIT RATIO	0.93	0.93	0	0.00%
PRIVOPs				
PRIVOP/CMD	18.735	19.154	0.419	2.24%
DIAG/CMD	16.806	16.588	-0.218	-1.30%
DIAG 08/CMD	0.732	0.722	-0.010	-1.37%
DIAG 10/CMD	5.565	5.493	-0.072	-1.29%
DIAG 58/CMD	1.227	1.213	-0.014	-1.14%
DIAG A4/CMD	3.997	3.945	-0.052	-1.30%
DIAG A8/CMD	1.862	1.890	0.028	1.50%
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Table 12. VM/ESA 1.0 Starter set: unmodified HCPIUF vs. HCPIUF with VMREAD change.

3.7 Migration

3.7.1 Introduction

In November and December of 1989, as the system was being developed, early measurements of the VM/ESA 1.0 system were obtained.

These measurements show the change in performance of the system as one migrates from VM/XA SP 2.0 or from VM/XA SP 2.1 to VM/ESA 1.0. Each component, CP, CMS, GCS and VTAM, was measured to see the performance change of that component. This was done by taking the existing PID level VM/XA SP 2.0 and VM/XA SP 2.1 systems and changing each VM/XA SP component to the same ESA 1.0 component. The order in which the components were changed was CP, GCS, CMS and VTAM. This provided an opportunity to isolate the performance change of each of those components.

Since these measurements were done on early levels of the VM/ESA 1.0 components, the results individual customers may see, as they install and migrate to the VM/ESA 1.0 system, may be different.

3.7.2 CMS Intensive Measurements

3.7.2.1 3090-600J / 512M / VM/XA SP 2.0 and VM/XA SP 2.1

Following is the description of the environment used to do the VM/XA SP 2.0 and VM/XA SP 2.1 to VM/ESA 1.0 Migration measurements.

```
1) WORKLOAD:     FS7B0R

2) HARDWARE CONFIGURATION

   - PROCESSOR:  3090-600J
   - STORAGE:
      - RSTOR:   512M
      - XSTOR:   2G

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PSYS02     3380      5
                 WKLD01     3380     16
                 WKLD02     3380     18

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - PAGE:          20         3380            3
      - SPOOL:          8         3380            3
      - TDISK:         12         3380            3
      - USER:          80         3380            7

   - COMMUNICATIONS:  CONTROLLER  NUMBER  LINES  LINESPEED
                      ----------  ------  -----  ---------
                        3725         2     128     19.2K
                        3745         1      50     56K

3) SOFTWARE CONFIGURATION

   - DRIVER:           TPNS
   - USER THINK TIME:  BACTRIAN (Avg 30 Secs)
   - CMS BLOCK SIZE:   4K
   - USER VM SIZE:     2M
   - USER RELSHARE:    100

                       TYPE       NUMBER  VM SIZE  RELSHARE  OTHER OPTIONS
                       ----       ------  -------  --------  -------------
   - SERVER MACHINE:   VTAM/VSCS    3       64M      10000    QUICKDSP ON
                       EXT. VSCS    3       64M      10000    QUICKDSP ON

4) MEASUREMENT SUMMARY
```

The following table summarizes the results of the early measurements where migration from VM/XA SP 2.0 and from VM/XA SP 2.1 to VM/ESA 1.0 was made. These measurements were done on early levels of all VM/ESA 1.0 components. The first column contains the results of the measurements of VM/XA SP 2.0 and the second column the results of VM/XA SP 2.1. The remaining columns give the results where an VM/ESA 1.0 component has replaced an XA SP component; CP, GCS, CMS and VTAM respectively. Each column has one additional VM/ESA 1.0 component. The last column contains not only the results of the change from VTAM 3.2 to VTAM 3.3 but the complete VM/ESA 1.0 system as well. As can be seen from looking at the key indicators (response time, throughput and processor utilization), VM/ESA 1.0 has a performance improvement over both VM/XA SP systems.

	XA SP2.0	XA SP2.1	CP/ESA	GCS/XA	CMS/ESA	VTAM 3.3
Storage						
REAL	512M	512M	512M	512M	512M	512M
XSTOR	2048M	2048M	2048M	2048M	2048M	2048M
Run ID	Y62R5201	Y6$R5200	Y63R5203	Y63R5204	Y63R5205	Y63R5206
Variables						
USERS	5200	5200	5200	5200	5200	5200
VTAMs	3	3	3	3	3	3
VSCSs	6	6	6	6	6	6
PROCESSORS	6	6	6	6	6	6
Response Time						
TRIV INT	0.068	0.068	0.074	0.073	0.066	0.067
NONTRIV INT	0.374	0.374	0.409	0.392	0.346	0.340
TOT INT	0.225	0.225	0.242	0.233	0.213	0.210
AVG FIRST (T)	0.806	0.806	0.919	0.886	0.703	0.703
AVG LAST (T)	1.049	1.046	1.189	1.143	0.913	0.913
Throughput						
ETR (T)	179.25	179.42	179.29	179.04	181.16	180.58
ITR (H)	33.75	34.00	33.41	33.38	34.76	34.49
ITR	42.58	42.93	43.28	42.99	42.92	42.64
EMUL ITR	66.75	67.39	70.16	69.21	69.73	69.27
ITRR (H)	1.000	1.007	0.990	0.989	1.030	1.022
ITRR	1.000	1.008	1.016	1.010	1.008	1.002
Proc Usage						
PBT/CMD (H)	29.629	29.415	29.927	29.962	28.769	28.993
CP/CMD (H)	11.469	11.290	11.950	11.865	11.478	11.571
EMUL/CMD (H)	18.157	18.122	17.974	18.094	17.289	17.420
Processor Util						
TOTAL (H)	531.10	527.75	536.54	536.43	521.18	523.55
UTIL/PROC (H)	88.52	87.96	89.42	89.41	86.86	87.26
TOTEMUL (H)	325.46	325.15	322.25	323.95	313.18	314.55
TVR (H)	1.63	1.62	1.66	1.66	1.66	1.66
Storage						
WKSET (V)	na	53	52	53	55	56
PGBLPGS	97586	97287	107K	107K	107K	107K
SHRPGS	986	1033	47K	46K	48K	54K
Paging						
READS/SEC	118	118	128	128	137	146
WRITES/SEC	6	8	22	28	43	58
PAGE/CMD	0.692	0.702	0.837	0.871	0.994	1.130
XSTOR IN/SEC	989	970	1066	1021	1099	1088
XSTOR OUT/SEC	1051	1043	1153	1113	1189	1189
XSTOR/CMD	11.381	11.220	12.377	11.919	12.630	12.609
I/O						
VIO RATE	1904	1923	1900	1905	1801	1789
VIO/CMD	8.433	8.509	8.191	8.261	8.078	8.021
MDC READS	1146	1149	1137	1143	1065	1069
MDC WRITES	486	487	484	487	505	501
MDC MODS	401	403	400	403	411	409
MDC HIT RATIO	0.93	0.93	0.93	0.93	0.91	0.92
PRIVOPs						
PRIVOP/CMD	21.203	23.497	23.229	22.885	25.249	25.300
DIAG/CMD	18.448	18.880	18.423	18.057	18.126	18.071
DIAG 08/CMD	0.736	0.736	0.725	0.726	0.729	0.731
DIAG 10/CMD	5.049	4.994	5.042	5.016	5.388	5.410
DIAG 58/CMD	1.222	1.221	1.216	1.218	1.220	1.224
DIAG A4/CMD	4.815	4.810	4.791	4.809	4.068	4.081
DIAG A8/CMD	3.269	3.383	3.285	3.295	3.356	3.295
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM						

Table 13. VM/XA SP 2.0 and VM/XA SP 2.1 Migration to VM/ESA 1.0

3.8 Single VTAM

3.8.1 Introduction

The measurements and analysis for single VTAM were completed in order to compare performance to multiple VTAM configurations, and to determine limitations in the single VTAM environment. Refer to Table 14 on page 74 for details on measurements with VTAM and VSCS in the same virtual machine. Table 15 on page 77 contains information on measurements with a single VTAM and 2 separate VSCS virtual machines. In addition, data on the comparison of a single VTAM configuration to one using three copies of VTAM with three internal and three external VSCSs is shown.

3.8.2 Measurements

3.8.2.1 3090-600J / 512M / Single VTAM VSCS Virtual Machine

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

- PROCESSOR: 3090-600J
- STORAGE:
  - RSTOR: 512M
  - XSTOR: 2G

- DASD:

PACK NAME	TYPE	CHANNEL
PSYS02	3380	5
WKLD01	3380	16
WKLD02	3380	18

	NUMBER OF PACKS	TYPE	NUMBER OF CHANNELS
- SYSTEM:			
- PAGE:	20	3380	3
- SPOOL:	8	3380	3
- TDISK:	12	3380	3
- USER:	80	3380	7

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES	LINESPEED
3725	2	128	19.2K
3745	1	50	56K

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- USER THINK TIME: BACTRIAN (Avg 30 Secs)
- CMS BLOCK SIZE: 4K
- USER VM SIZE: 2M
- USER RELSHARE: 100

	TYPE	NUMBER	VM SIZE	RELSHARE	OTHER OPTIONS
- SERVER MACHINE:	VTAM/VSCS	1	64M	10000	QUICKDSP ON

4) MEASUREMENT SUMMARY

This set of measurements reflects the capacity of a single VTAM/VSCS virtual machine to support remote users with the FS7B0R workload. The 3900 user and 4200 user measurements are easily supported by the system, with processor utilization of 59% and 65% respectively. TPNS response

time was subsecond in both measurements. Addition of a VSCS machine would be indicated by excessively high CPU resource consumption. Measurements indicate that 4500 users can not be supported on this workload for this environment. Full data wasn't collected for the 4500 user measurement. The data supplied indicates a considerable increase in response time rates, along with a decrease in overall processor utilization. Thus,

- the capacity of a single VTAM/VSCS virtual machine is between 4200 and 4500 users.
- the limiting factor using this workload is not storage as it was when using VTAM 3.2.

	ESA 1.0	ESA 1.0	ESA 1.0
Storage			
REAL	512M	512M	512M
XSTOR	2048M	2048M	2048M
Run ID	Y63R3902	Y63R4201	Y63R4500
Variables			
USERS	3900	4200	4500
VTAMs	1	1	1
VSCSs	1	1	1
PROCESSORS	6	6	6
Response Time			
TRIV INT	0.035	0.043	0.085
NONTRIV INT	0.249	0.287	5.428
TOT INT	0.161	0.182	1.956
AVG FIRST (T)	0.413	0.420	na
AVG LAST (T)	0.577	0.583	na
Throughput			
ETR (T)	138.56	149.34	na
ITR (H)	39.19	38.48	na
ITR	44.44	44.98	na
EMUL ITR	75.09	77.26	na
ITRR (H)	1.000	0.982	na
ITRR	1.000	1.012	na
Proc Usage			
PBT/CMD (H)	25.518	25.984	na
CP/CMD (H)	10.820	11.234	na
EMUL/CMD (H)	14.695	14.748	na
Processor Util			
TOTAL (H)	353.57	388.04	286
UTIL/PROC (H)	58.93	64.67	47.67
TOTEMUL (H)	203.62	220.25	na
TVR (H)	1.74	1.76	na
Storage			
WKSET (V)	62	60	na
PGBLPGS	111K	110K	na
SHRPGS	1010	1019	na
Paging			
READS/SEC	140	135	na
WRITES/SEC	29	33	na
PAGE/CMD	1.220	1.125	na
XSTOR IN/SEC	698	797	na
XSTOR OUT/SEC	764	865	na
XSTOR/CMD	10.552	11.129	na
I/O			
VIO RATE	1339	1462	na
VIO/CMD	8.535	8.403	na
MDC READS	773	834	na
MDC WRITES	349	385	na
MDC MODS	300	330	na
MDC HIT RATIO	0.94	0.93	na
PRIVOPs			
PRIVOP/CMD	20.751	20.438	na
DIAG/CMD	17.591	17.356	na
DIAG 08/CMD	0.729	0.737	na
DIAG 10/CMD	5.131	5.123	na
DIAG 58/CMD	1.241	1.239	na
DIAG A4/CMD	3.811	3.830	na
DIAG A8/CMD	3.212	3.321	na

	ESA 1.0	ESA 1.0	ESA 1.0
Storage			
REAL	512M	512M	512M
XSTOR	2048M	2048M	2048M
Run ID	Y63R3902	Y63R4201	Y63R4500
Variables			
USERS	3900	4200	4500
VTAMs	1	1	1
VSCSs	1	1	1
PROCESSORS	6	6	6
VTAM & VSCS Machines			
WKSET (V)	2285	2657	na
TOT CPU/CMD (V)	4.723	4.757	na
CP CPU/CMD (V)	2.951	3.025	na
EMUL/CMD (V)	1.772	1.732	na
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM			

Table 14. Single VTAM/VSCS Virtual Machine Measurements

3.8.2.2 3090-600J / 512M / VTAM, 2 Separate VSCS Virtual Machines

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

- PROCESSOR: 3090-600J
- STORAGE:
  - RSTOR: 512M
  - XSTOR: 2G

- DASD:

	PACK NAME	TYPE	CHANNEL
- SYSTEM:	PSYS02	3380	5
	WKLD01	3380	16
	WKLD02	3380	18

	NUMBER OF PACKS	TYPE	NUMBER OF CHANNELS
- PAGE:	20	3380	3
- SPOOL:	8	3380	3
- TDISK:	12	3380	3
- USER:	80	3380	7

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES	LINESPEED
3725	2	128	19.2K
3745	1	50	56K

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- USER THINK TIME: BACTRIAN (Avg 30 Secs or 45 Secs as indicated)
- CMS BLOCK SIZE: 4K
- USER VM SIZE: 2M
- USER RELSHARE: 100

	TYPE	NUMBER	VM SIZE	RELSHARE	OTHER OPTIONS
- SERVER MACHINE:	VTAM	1	64M	10000	QUICKDSP ON
	EXT. VSCS	2	64M	10000	QUICKDSP ON

4) MEASUREMENT SUMMARY

This set of measurements compares a single VTAM configuration with two external VSCS virtual machines to a configuration containing three VTAMs with three internal and three external VSCSs. In addition, the 7760 user measurement reflects the upper range in number of users supported on the FS7B0R workload for this single VTAM configuration. The CPU percentage, ITR, and ETR indicators for the two 5700 user measurements which compare a single VTAM to a three VTAM configuration are all within 3% of each other. Notice, however, that the response time as measured by TPNS has degraded somewhat. This can be attributed to a combination of two factors. First, in the three VTAM configuration, there are a total of six separate virtual machines (three VTAM/VSCS, and three VSCS) which may be separately scheduled on the six processors of the 3090 to overlap processing for the 5700 users. However, in the single VTAM configuration, it is possible to schedule overlapping processing for the same number of users on only three processors.

The second factor affecting response time in the 5700 user measurements is the fact that, in the single VTAM configuration, all VSCS work is in a separate virtual machine from VTAM. This creates extra pathlength for both VTAM and GCS due to the IUCV communications required to communicate between the VTAM and VSCS virtual machines. In the three VTAM configuration, each VTAM virtual machine also contained VSCS. Therefore, for the users logged on to these three VSCSs, the extra IUCV overhead was not incurred, and did not add to response time.

In the 7760 user measurement, the think time value was raised from the average of 30 seconds used in the 5700 user measurements. An average of 45 seconds was used in the 7760 user measurement.

Notice that the system was quite busy in the 5700 user measurements (processor utilization of 88.8 and 87.3% for the two measurements) with a 30 second think time. In order to support the additional users without overloading the system, the longer think time was used.

	ESA 1.0	ESA 1.0	ESA 1.0
Storage			
REAL	512M	512M	512M
XSTOR	2048M	2048M	2048M
Run ID	Y63R5702	Y63R5701	Y63R7761
Variables			
USERS	5700	5700	7760
VTAMs	3	1	1
VSCSs	6	2	2
PROCESSORS	6	6	6
Response Time			
TRIV INT	0.067	0.065	0.043
NONTRIV INT	0.399	0.514	0.374
TOT INT	0.246	0.278	0.214
AVG FIRST (T)	0.804	1.128	0.940
AVG LAST (T)	1.046	1.439	1.192
Throughput			
ETR (T)	196.10	192.83	179.36
ITR (H)	37.37	36.19	35.20
ITR	46.61	51.82	46.15
EMUL ITR	79.58	91.10	82.37
ITRR (H)	1.000	0.968	0.942
ITRR	1.000	1.112	0.990
Proc Usage			
PBT/CMD (H)	26.760	27.632	28.407
CP/CMD (H)	11.565	12.411	13.014
EMUL/CMD (H)	15.192	15.218	15.390
Processor Util			
TOTAL (H)	524.75	532.82	509.51
UTIL/PROC (H)	87.46	88.80	84.92
TOTEMUL (H)	297.90	293.43	276.05
TVR (H)	1.76	1.82	1.85
Storage			
WKSET (V)	55	53	na
PGBLPGS	104K	104K	96230
SHRPGS	1526	1152	1263
Paging			
READS/SEC	206	205	500
WRITES/SEC	117	127	368
PAGE/CMD	1.647	1.722	4.840
XSTOR IN/SEC	1292	1450	1319
XSTOR OUT/SEC	1461	1624	1760
XSTOR/CMD	14.039	15.942	17.167
I/O			
VIO RATE	1913	1862	1750
VIO/CMD	7.825	6.743	7.457
MDC READS	1104	1089	1016
MDC WRITES	526	521	485
MDC MODS	435	431	394
MDC HIT RATIO	0.92	0.92	0.91
PRIVOPs			
PRIVOP/CMD	34.511	36.210	24.631
DIAG/CMD	26.330	26.226	17.506
DIAG 08/CMD	0.724	0.731	0.714
DIAG 10/CMD	5.186	5.191	5.207
DIAG 58/CMD	1.234	1.219	1.232
DIAG A4/CMD	3.906	3.915	3.914
DIAG A8/CMD	3.315	3.226	3.317

	ESA 1.0	ESA 1.0	ESA 1.0
Storage			
REAL	512M	512M	512M
XSTOR	2048M	2048M	2048M
Run ID	Y63R5702	Y63R5701	Y63R7761
Variables			
USERS	5700	5700	7760
VTAMs	3	1	1
VSCSs	6	2	2
PROCESSORS	6	6	6
VTAM & VSCS Machines			
WKSET (V)	4136	3617	na
TOT CPU/CMD (V)	4.754	7.150	na
CP CPU/CMD (V)	2.881	5.197	na
EMUL/CMD (V)	1.873	1.953	na
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM			

Table 15. Single VTAM with Separate VSCS Virtual Machine Measurements

3.9 MVS Guest

3.9.1 Introduction

MVS Guest measurements were taken under the VM/XA SP 2.1 and VM/ESA 1.0 systems. The results were compared with MVS native performance to ascertain that there was no guest performance degradation. The following are the details of the measurements.

3.9.2 MVS Guest Measurements

The following is the run description for the MVS Guest measurements.

```
1) WORKLOAD:     CB84

2) HARDWARE CONFIGURATION

   - PROCESSOR:  3090-600J   CPU 1 only
   - STORAGE:
      - RSTOR:   256M
      - XSTOR:   512M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PSYS02     3380      5
                 DRV308     3380      6
                 USERPK     3380      6
                 PPLOAD     3380      8
                 PAGEPK     3380     11
                 PAGEF1     3380     14
                 PROGPK     3380     14
                 POOLS1     3380     16
                 POOLS2-6   3380      1
                 POOLS7     3380      7
                 POOLS8     3380      5
                 POOLS9     3380     15
                 CB8413     3380      6
                 CB84LB     3380     13
                 CBLOAD     3380      7
                 CKPTDS     3380      8
                 STGF31-37  3380     16
                 STG637     3380      6
                 IDAVOL     3380      6
                 PPL664     3380      6
                 PPL665     3380      6


      - CB84 USE:  NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                   ---------------  ----  ------------------
                         16         3350           2
                        160         3380          10

   - TAPE:       MONITOR 3480

3) SOFTWARE CONFIGURATION

   - DRIVER:           MA03VR12
   - VM SIZE:          256M
   - V=R SIZE:         212M
   - V=F SIZE:         212M
   - GUEST PRIORITY:   1
```

4) MEASUREMENT SUMMARY

The following table contains the measurement data for the CB84 runs that were made for ESA 1.0 For comparative purposes, columns with MVS native data and XA 2.1 data are included with the ESA 1.0 data. RMF and hardware monitor are the sources for the measurement data.

Guest operating system performance on ESA 1.0 was predicted to be equivalent to XA 2.1 because all of the guest performance improvements made to XA 2.0 and XA 2.1 were propagated in ESA 1.0. The results of these measurements, as detailed in the following table, show that to be the case.

	MVS 310e	XA 2.1	ESA 1.0	ESA 1.0	ESA 1.0
Run ID	Y1YY24B2	Y1$Y23B1	Y13Y0004	Y13Y0005	Y13Y0006
Guest Type	Native	V = R	V = R	V = F	V = V
Configuration					
Processor Type	3090J	3090J	3090J	3090J	3090J
Num. Processors	1	1	1	1	1
Real Storage	256M	256M	256M	256M	256M
Expanded Storage	512M	512M	512M	512M	512M
MVS Sys Lvl	3.1.0e	3.1.0e	3.1.0e	3.1.0e	3.1.0e
Throughput					
Int Throughput (ITR)	0.473	0.450	0.453	0.453	0.309
ITR % of Native	100	95.1	95.8	95.8	65.3
Ext Throughput (ETR)	0.461	0.440	0.442	0.441	0.303
ETR % of Native	100	95.4	95.9	95.7	65.7
Processor Data					
Elapsed Seconds	846.290	886.516	883.309	883.469	1285.330
Processor Busy %	97.51	97.81	97.54	97.49	98.35
Processor Seconds	825.220	867.101	861.580	861.290	1264.122
Batch Data					
Num. of Initiators	24	23	23	23	22
Num. of Batch Jobs	390	390	390	390	390
Chan Path/DASD Data					
No. of DASD Paths	32	32	32	32	32
Avg. Ch Path Busy	8.68	9.10	8.82	8.96	6.65
High Ch Path Busy	23.42	24.67	25.17	25.17	18.79
I/O Interrupt Rate	588.80	559.00	544.30	544.10	378.20
I/O Interrupts/Tran	1277.68	1270.67	1232.80	1232.55	1248.48
Paging					
Total: In+Out	0.00	0.00	0.00	0.00	0.00
In	0.00	0.00	0.00	0.00	0.00
NSW/NVIO: Total	0.00	0.00	0.00	0.00	0.00
In	0.00	0.00	0.00	0.00	0.00
VIO: Total	0.00	0.00	0.00	0.00	0.00
In	0.00	0.00	0.00	0.00	0.00
Main Storage Data					
Storage Size	127M	128M	128M	128M	128M
Total Frames					
LPA - Avg	891	775	695	691	696
CSA - Avg	282	279	278	278	277
Priv Area - Avg	8140	7151	7011	7002	7016
Unused - Avg	15780	21743	21968	21976	21976
Total - Avg	32318	32766	32766	32766	32766
Fixed Frames					
SQA - Avg	4421	819	813	814	813
LPA - Avg	33	32	33	33	33
CSA - Avg	0	0	0	0	0
LSQA - Avg	966	939	942	944	929
Priv Area - Avg	984	886	873	887	895
Below 16M - Avg	308	270	279	278	273
Tot Fixed - Avg	8093	3586	3567	3586	3588
Nucleus Frames	1675	907	907	907	907

Table 16. CB84 Measurement Data: ESA 1.0 And XA 2.1 MVS/3.1.0e Guests

4.0 Specific Measurements: VM/ESA 1.0 370 Feature

4.1 Introduction

This chapter contains the configuration details and specific results for those measurements obtained to evaluate the performance of VM/ESA 1.0 370 Feature.

4.1.1 Format Description

For each group of measurements in this section there are two parts. The first part describes the run environment and data while the second presents the data in table format. The measurement descriptions have four sections. These are:

1. Workload: Gives the name of the workload associated with the data. For more detail see section "Appendix C. Workloads" on page 163.
2. Hardware Configuration: This section gives a high-level overview of the hardware configuration. It contains the following subsections:
   - Processor: The processor on which the data was collected.
   - Storage: The amount of real storage available on the processor.
   - System DASD: The packs on which the system resides.
   - User/Page DASD: Page space is spread across multiple user packs. This represents the number of user packs on the system and the number of channels they're spread across. This also includes spool and TDISK space. For more information on the pack layout see section "Appendix D. Configuration Details" on page 182.
3. Software Configuration: This section contains pertinent software information. It breaks down into the following subsections:
   - Driver: The tool used to run the workload.
   - User Think Time: Average amount of time between user interactions. See the following discussion.
   - CMS Blocksize: The blocksize of CMS minidisks.
   - User VM Size: The storage size of the user virtual machine.
   - User Priority: The users' scheduling priority.
   - Server Machines: Lists the type of server machine, how many there are, their storage size, scheduling priority, and any special options used.
4. Measurement Summary: This contains a brief overview of the measurement data.

Think Times Measurements in this document use several different think times. All of them stem from two distributions - the bactrian curve and the uniform curve. The CMS interactive workloads on the high end use the bactrian think time, while the 370 Feature measurements run with the uniform think time. Both have a basis in actual user studies. Low end measurements are unable to attain a large enough sample size to insure repeatability in a practical measurement period using the bactrian distribution. For this reason, 370 Feature measurements in this document use uniformly distributed random think times which represent only active users. Note that user think times denoted as random within this section are uniformly distributed.

In the 370 Feature measurements there are also two different average think times based on the uniform distribution. These are a 5 second and a 10 second average think time. Both the Connectivity and some CRR measurements use the 5 second think times. This allows the Connectivity measurements to achieve the desired CPU utilization with a relatively low number of users. It also allows the CRR measurements to maintain the desired syncpoint rates. Measurements using the CMS interactive workload use the 10 second average think time. This average is based on actual user data.

Refer to the glossary for additional definitions of performance terms used in this chapter.

4.1.2 Tools Description

The data in this section comes from several tools. These include VMMAP, RESPONS2, TPNS, FSID, and QUERY FILEPOOL STATUS. VMMAP and RESPONS2 both reduce VM/370 MONITOR data. VMMAP (Virtual Machine Facility / 370 Performance / Monitor Analysis) is a program product which produces reports summarizing the performance and utilization of the system. RESPONS2 is an internally developed tool which reduces RESPONSE and SCHEDULE class MONITOR records, producing a report on response times and resource utilization for the commands found in the MONITOR data. TPNS (Teleprocessing Network Simulator) is a program product terminal and network simulation tool which provides system performance and response time information. Full Screen Internal Driver (FSID) is an internally developed tool that simulates locally attached full-screen terminals (3277′s). QUERY FILEPOOL STATUS is a CMS command that provides information on activity within SFS filepool servers and the CRR recovery server.

4.1.3 Additional Considerations

Tables in this section include the Internal Throughput Rate (ITR) as a measure of throughput. The ITR is the number of units of work accomplished per unit of CPU busy time in an unconstrained environment. Some of the measurements in this section have constraints in paging, making their ITRs not fit the definition exactly. Since there are constraints in these environments, these ITRs are not valid for comparisons across processors although they are still a useful measure of throughput. These have been marked with an * to highlight the system constraints.

4.2 CMS Regression

4.2.1 Introduction

The purpose of this section is to evaluate performance differences that have occurred due to the addition of new function and service since the introduction of VM/SP Release 6. The workload executed has not changed to explicitly use any of the new ESA 1.0 functions. Therefore, the same end user work was completed in both cases. These comparisons can be used as an indicator of what might be expected with similar workloads when up-leveling to this new level of VM.

SFS counts and timings obtained from the QUERY FILEPOOL STATUS command are provided in “Appendix B. SFS Counter Data” on page 155 to supplement the SFS measurement data provided in this section.

4.2.2 Measurements

4.2.2.1 4381-13 / 16M / 0% SFS

```
1) WORKLOAD:     FS7B0

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380     01
                 PRF01      3380     02
                 PRFPRO     3380     03

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      10        3380            4


3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64


4) MEASUREMENT SUMMARY
```

This comparison does not include any Shared File System activity, only the EDF File System was used. The data shows a reduction in storage requirements, a large reduction in CPU usage, and a slight reduction in virtual file I/Os.

The storage required by CMS has changed in the following ways: The CMS shared nucleus area has grown, the non-shared low nucleus area has grown, and the storage referenced for moderate to large accessed disks has been reduced. The net effect of these changes is very workload dependent. The workload and environment measured here netted a reduction in the overall paging rate of 10% at the high end of system activity.

The Processor Busy Time (PBT) used by CMS has also changed. Several individual changes were made, with the most notable ones being reduced processor busy time for EXECIO when called from REXX, COPYFILE improvements on a per record basis, Assembler time reductions due to an improvement in CMSSTOR macro expansion, and an increase in the IPL CMS time. The

EXECIO improvement is also responsible for the large reduction in the PRIVOP rate. SPKA instructions were eliminated from the linkage path between the REXX and EXECIO modules.

The virtual I/O rate showed a slight reduction, mainly due to a reduction in I/Os for the Assembler when expanding the CMSSTOR macro.

The combination of the above performance changes netted a reduction in average response time (AVG TOT (R)) of 2% to 20%.

	VM/SP 6			VM/ESA 1.0		
Real Storage						
TOTAL STORAGE	16384K			16384K		
NUCLEUS	463K			488K		
DYNAMIC	14292K			14244K		
FREE	1392K			1395K		
TRACE	236K			256K		
Run-ID	EC4295	EC4295	EC4295	EC7603	EC7603	EC7603
ACTIVE USERS	69	77	85	69	77	85
Response Time						
TRIV RESP	0.23	0.31	0.43	0.23	0.29	0.38
1ST RESP (R)	0.44	0.63	0.89	0.42	0.55	0.70
AVG TOT (R)	0.63	0.86	1.22	0.62	0.80	0.98
Throughput						
ITR	9.14	8.75	* 8.35	9.77	9.41	* 9.05
ETR (R)	6.348	7.023	7.534	6.431	7.121	7.643
CPU Usage						
PBT/CMD	0.109	0.114	0.120	0.102	0.106	0.111
CP/CMD	0.026	0.028	0.032	0.026	0.029	0.031
VIRT/CMD	0.084	0.086	0.088	0.076	0.078	0.079
PBT/CMD (R)	0.104	0.107	0.112	0.097	0.099	0.102
CP/CMD (R)	0.020	0.022	0.024	0.021	0.022	0.023
VIRT/CMD (R)	0.084	0.085	0.088	0.076	0.077	0.079
CPU Utilization						
TOTCPU	69.46	80.28	90.22	65.85	75.73	84.44
CPCPU	16.17	19.97	23.89	16.82	20.28	23.73
VIRTCPU	53.29	60.31	66.33	49.03	55.45	60.71
Storage						
WKSET	205.13	193.36	179.00	204.18	189.50	178.96
PGBLPGS	3565.00	3565.00	3565.00	3553.00	3553.00	3553.00
SHRPGS	309.84	340.50	319.16	312.21	327.42	324.89
Paging						
PAGE RATE	96.48	140.55	187.04	95.96	129.57	167.85
PAGE/CMD	15.20	20.01	24.83	14.92	18.20	21.96
I/0						
VIO RATE	67.00	74.08	80.13	65.42	72.62	78.40
VIO/CMD	10.55	10.55	10.64	10.17	10.20	10.26
MDSK/CMD (R)	5.002	4.956	5.078	4.700	4.676	4.712
Queueing						
PCTCPUQ	1.98	3.62	4.26	2.73	2.52	2.96
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	1.75	3.04	3.45	2.05	2.11	2.84
PCTIOQ	1.60	1.10	1.23	1.67	1.36	1.85
PRIVOPs						
PRIVOP RATE	161.51	181.85	202.61	68.39	74.83	79.67
PRIVOP/CMD	25.44	25.89	26.89	10.63	10.51	10.42
Note: R=RESPONS2, Unmarked=VMMAP, * = not a true ITR due to system constraints.						

Table 17. 4381-13 / 16M / 0% SFS

> *4.2.2.2 4381-13 / 16M / 0% SFS - SP5 Comparisons*

```
> 1) WORKLOAD:     FS7B0

> 2) HARDWARE CONFIGURATION

>    - PROCESSOR:  4381-13
>    - STORAGE:    16M

>    - DASD:       PACK NAME  TYPE  CHANNEL
>                  ---------  ----  -------
>       - SYSTEM:  PRFRES     3380     01
>                  PRF01      3380     02
>                  PRFPRO     3380     03

>                  NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
>                  ---------------  ----  ------------------
>       - USER/PAGE:      10        3380            4


> 3) SOFTWARE CONFIGURATION

>    - DRIVER:           FSID
>    - USER THINK TIME:  RANDOM (Avg 10 Secs)
>    - CMS BLOCKSIZE:    4K
>    - USER VM SIZE:     2M
>    - USER PRIORITY:    64


> 4) MEASUREMENT SUMMARY
```

> The following comparison is included to demonstrate the impact of migrating from VM/SP5, service level 500 with CMS at service level 508, to ESA 1.0. The deltas in the comparison can be attributed to many factors wich have been introduced through new functions and service in CMS 5.5, VM/SP 6, and ESA 1.0.

> The CP nucleus and trace table have grown since VM/SP 5. The number of SHRPGS has drastically grown due to CMS growth and the addition of CSL (Callable Services Library). The page rate has grown because of the growth in CP, CMS and also the non-shared page growth. There has been a slight decrease in VIOs which was due to the write-behind function introduced in VM/SP 6. EXECIO improvement is responsible for the large reduction in the PRIVOP rate. SPKA instructions were eliminated from the linkage path between the REXX and EXECIO modules. Segment Protect and 4K key support became available through service to VM/SP 5 and this resulted in a 6% improvement in ITR. These changes resulted in the average first response time increasing by 0.26 seconds, which would meet our definition of equivalence; and an ITR decrease of 13.5% which would not meet our definition of equivalence.

> The individual commands which have grown the most since VM/SP 5 are Assembler, HASM, IPL, productivity aids, and Xedit. The Assembler and HASM changes are due to CMS 5.5 which introduced CMSSTOR. IPL grew because of the CMS growth. The productivity aids growth was introduced in VM/SP 6 with EXEC2 changes to REXX for some of their macros. Xedit growth is due to changes needed for SFS in VM/SP 6.

> A considerable amount of new function was introduced since VM/SP 5. Most of these deltas are related to the associated storage growth and can be remedied by the >16 Meg support provided by ESA 1.0. See section “Greater Than 16M of Real Storage” on page 135 for more details.

	VM/SP 5			VM/ESA 1.0		
Real Storage						
TOTAL STORAGE		16384K			16384K	
NUCLEUS		463K			488K	
DYNAMIC		14292K			14244K	
FREE		1392K			1395K	
TRACE		236K			256K	
Run-ID	EC7620	EC7620	EC7620	EC7603	EC7603	EC7603
ACTIVE USERS	69	77	85	69	77	85
Response Time						
TRIV RESP	0.17	0.21	0.26	0.23	0.29	0.38
1ST RESP (R)	0.30	0.37	0.44	0.42	0.55	0.70
AVG TOT (R)	0.49	0.58	0.68	0.62	0.80	0.98
Throughput						
ITR	10.95	10.79	10.46	9.77	9.41	* 9.05
ETR (R)	6.472	7.168	7.883	6.431	7.121	7.643
CPU Usage						
PBT/CMD	0.091	0.093	0.096	0.102	0.106	0.111
CP/CMD	0.024	0.026	0.027	0.026	0.029	0.031
VIRT/CMD	0.067	0.067	0.068	0.076	0.078	0.079
PBT/CMD (R)	0.087	0.088	0.089	0.097	0.099	0.102
CP/CMD (R)	0.020	0.021	0.022	0.021	0.022	0.023
VIRT/CMD (R)	0.067	0.067	0.067	0.076	0.077	0.079
CPU Utilization						
TOTCPU	59.09	66.41	75.36	65.85	75.73	84.44
CPCPU	15.53	18.25	21.47	16.82	20.28	23.73
VIRTCPU	43.56	48.17	53.89	49.03	55.45	60.71
Storage						
WKSET	201.73	182.03	173.24	204.18	189.50	178.96
PGBLPGS	3544.00	3544.00	3544.00	3553.00	3553.00	3553.00
SHRPGS	235.26	240.42	240.95	312.21	327.42	324.89
Paging						
PAGE RATE	68.71	88.62	114.27	95.96	129.57	167.85
PAGE/CMD	10.62	12.36	14.50	14.92	18.20	21.96
I/0						
VIO RATE	68.00	74.99	82.54	65.42	72.62	78.40
VIO/CMD	10.51	10.46	10.47	10.17	10.20	10.26
MDSK/CMD (R)	5.000	4.936	4.967	4.700	4.676	4.712
Queueing						
PCTCPUQ	1.45	0.89	1.73	2.73	2.52	2.96
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	1.91	2.19	1.43	2.05	2.11	2.84
PCTIOQ	1.84	1.44	0.99	1.67	1.36	1.85
PRIVOPs						
PRIVOP RATE	156.17	169.72	185.50	68.39	74.83	79.67
PRIVOP/CMD	24.13	23.61	23.53	10.63	10.51	10.42
Note: R=RESPONS2, Unmarked=VMMAP, * = not a true ITR due to system constraints.						

Table 18. 4381-13 / 16M / 0% SFS

4.2.2.3 4381-13 / 16M / 35% SFS

```
1) WORKLOAD:     FS7B35

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380     01
                 PRF05      3380     01
                 PRF01      3380     02
                 PRF06      3380     02
                 PRFPRO     3380     03

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       10      3380            4


3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64

   - SERVER MACHINES:
                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
            - SP 6:    SFS          1      16M         1    QDROP OFF USERS
            - ESA 1.0: SFS          1      16M         1    QDROP OFF USERS
                       CRR          1      16M        64

4) MEASUREMENT SUMMARY
```

This comparison includes a combination of Shared File System and EDF File System activity. This measurement differs from the 0% SFS measurement above in that all user files are now in SFS. Only the installation supplied files and system files (that is, all files that are read only) remain on minidisks. The resulting reduction in minidisk I/Os per command is approximately 35% and the activity that caused them is assumed by the Shared File System. For example, using the middle column in ESA 1.0, MDSK/CMD (R) went from 4.676 for the 0% SFS run to 3.005 in this measurement, which represents a 35% decrease.

The data shows equivalent real storage requirements, a reduction in processor usage, and a slight reduction in virtual file I/Os as compared to VM/SP Release 6.

The storage requirements (PAGE/CMD under Paging) changed by -3% to +2%. This does not show the improvement demonstrated for the 0% SFS runs. The SFS server storage requirements are similar between releases, so the storage impact due to SFS occurs in the user virtual machines. In SP 6, Coordinated Resource Recovery (CRR) did not exist. With ESA 1.0 370 Feature, SFS is participating in CRR. The majority of the storage impact is due to new CRR control blocks which contain information on how SFS is participating in CRR. These new data areas are being referenced.

The processor busy time (PBT/CMD under CPU Usage) decreased by 4% to 6%, but did not decrease as much as for the 0% SFS runs. Contributing factors included:

- new CRR processing.
- no net decrease in storage requirements
- increased server processing

File pool requests per command (FP REQ/CMD (Q)) dropped significantly by 28% to 30% because of a performance enhancement added to ESA 1.0 that eliminates most of the times the user virtual machines request File Status Table (FST) update information from the SFS server.

Virtual file I/Os per command (sum of MDSK/CMD (R) and IO/CMD (Q)) decreased by 3% to 6%.

Average response time (AVG TOT (R)) improved by 6% to 13%. The degree of improvement is less than for the 0% SFS runs because PAGE/CMD did not improve like it did in the 0% SFS runs and because PBT/CMD showed a smaller improvement than it did in the 0% SFS runs.

	VM/SP 6			VM/ESA 1.0		
Real Storage						
TOTAL STORAGE		16384K			16384K	
NUCLEUS		463K			488K	
DYNAMIC		14292K			14244K	
FREE		1392K			1395K	
TRACE		236K			256K	
Run-ID	EC4292	EC4292	EC4292	EC7600	EC7600	EC7600
ACTIVE USERS	59	67	75	59	67	75
Response Time						
TRIV RESP	0.31	0.40	0.57	0.28	0.36	0.47
1ST RESP (R)	0.53	0.72	0.99	0.49	0.65	0.83
AVG TOT (R)	0.75	0.97	1.31	0.70	0.89	1.14
Throughput						
ITR	7.81	7.57	* 7.31	8.18	8.06	* 7.74
ETR (R)	5.398	6.018	6.478	5.468	6.133	6.619
CPU Usage						
PBT/CMD	0.128	0.132	0.137	0.122	0.124	0.129
CP/CMD	0.033	0.036	0.039	0.034	0.036	0.039
VIRT/CMD	0.095	0.097	0.098	0.089	0.088	0.090
PBT/CMD (R)	0.107	0.111	0.114	0.101	0.102	0.105
CP/CMD (R)	0.021	0.024	0.025	0.022	0.023	0.025
VIRT/CMD (R)	0.086	0.087	0.089	0.079	0.079	0.080
CPU Utilization						
TOTCPU	69.10	79.48	88.60	66.85	76.11	85.55
CPCPU	17.59	21.34	25.23	18.40	21.98	25.89
VIRTCPU	51.51	58.14	63.37	48.45	54.13	59.66
Storage						
WKSET	228.43	210.21	196.88	220.98	207.66	192.83
PGBLPGS	3565.00	3565.00	3565.00	3553.00	3553.00	3553.00
SHRPGS	319.53	332.05	329.63	398.79	404.37	411.84
Paging						
PAGE RATE	87.17	125.24	169.97	90.39	127.37	167.47
PAGE/CMD	16.15	20.81	26.24	16.53	20.77	25.30
I/0						
VIO RATE	47.97	53.50	57.22	46.75	52.09	56.32
VIO/CMD	8.89	8.89	8.83	8.55	8.49	8.51
MDSK/CMD (R)	3.315	3.329	3.298	3.048	3.005	3.015
Queueing						
PCTCPUQ	3.16	3.79	4.00	2.41	2.36	3.21
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	2.11	1.78	3.31	1.55	2.89	3.96
PCTIOQ	0.88	0.85	0.69	0.86	0.53	0.89
PRIVOPs						
PRIVOP RATE	189.21	202.29	218.40	90.53	98.98	105.80
PRIVOP/CMD	35.05	33.61	33.71	16.56	16.14	15.98
SFS Server						
WKSET	1480	1501	1514	1458	1490	1473
PAGE/CMD	0.00	0.15	0.27	0.00	0.14	0.27
PBT/CMD	0.011	0.013	0.013	0.014	0.013	0.014
CP/CMD	0.005	0.006	0.005	0.004	0.004	0.005
VIRT/CMD	0.006	0.007	0.008	0.009	0.009	0.009
FP REQ/CMD (Q)	2.11	2.09	2.10	1.52	1.49	1.47
IO/CMD (Q)	2.36	2.43	2.50	2.43	2.39	2.41
IO TIME/CMD (Q)	0.07	0.08	0.11	0.07	0.08	0.10
SFS TIME/CMD (Q)	0.10	0.13	0.16	0.10	0.12	0.14

Note: R=RESPONS2, Q=QUERY FILEPOOL STATUS, Unmarked=VMMAP, * = not a true ITR due to system constraints.

Table 19. 4381-13 / 16M / 35% SFS

4.2.2.4 4381-13 / 16M / MAX% SFS

```
1) WORKLOAD:     FS7BMAX

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380     01
                 PRF05      3380     01
                 PRF01      3380     02
                 PRF06      3380     02
                 PRFPRO     3380     03

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      10        3380            4


3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64

   - SERVER MACHINES:
                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
            - SP 6:    SFS          1      16M         1    QDROP OFF USERS
            - ESA 1.0: SFS          1      16M         1    QDROP OFF USERS
                       CRR          1      16M        64

4) MEASUREMENT SUMMARY
```

This comparison has all files which SFS supports placed in SFS directories. This includes all user files and the installation supplied files. Only the S and Y disks, which SFS does not support, remain on minidisks. The resulting reduction in minidisk I/Os per command is approximately 48% (when compared to “4381-13 / 16M / 0% SFS” on page 84), and the activity that caused them is assumed by the Shared File System. For example, using the middle column in ESA 1.0, MDSK/CMD(R) went from 4.676 for the 0% SFS run to 2.444 in this measurement, which represents a 48% decrease.

The data shows an increase in storage requirements (as evidenced by a 3% to 10% increase in PAGE/CMD), a slight decrease in processor usage of 2% to 3%, and a slight decrease in virtual file I/Os per command of 3% to 4%. The response times dropped an average of 0.03 seconds.

	VM/SP 6			VM/ESA 1.0		
Real Storage						
TOTAL STORAGE		16384K			16384K	
NUCLEUS		463K			488K	
DYNAMIC		14292K			14244K	
FREE		1392K			1395K	
TRACE		236K			256K	
Run-ID	EC4599	EC4599	EC4599	EC7604	EC7604	EC7604
ACTIVE USERS	57	65	73	57	65	73
Response Time						
TRIV RESP	0.23	0.36	0.44	0.24	0.31	0.44
1ST RESP (R)	0.47	0.68	0.90	0.46	0.60	0.89
AVG TOT (R)	0.67	0.95	1.21	0.68	0.85	1.21
Throughput						
ITR	7.70	7.40	* 7.10	7.93	7.59	* 7.30
ETR (R)	5.239	5.826	6.449	5.273	5.934	6.387
CPU Usage						
PBT/CMD	0.130	0.135	0.141	0.126	0.132	0.137
CP/CMD	0.033	0.036	0.039	0.034	0.037	0.041
VIRT/CMD	0.097	0.100	0.102	0.092	0.095	0.097
PBT/CMD (R)	0.107	0.109	0.111	0.101	0.102	0.106
CP/CMD (R)	0.021	0.022	0.024	0.022	0.023	0.025
VIRT/CMD (R)	0.086	0.087	0.087	0.079	0.079	0.081
CPU Utilization						
TOTCPU	68.03	78.77	90.87	66.51	78.17	87.48
CPCPU	17.10	20.76	25.15	18.03	22.05	25.84
VIRTCPU	50.92	58.01	65.72	48.48	56.12	61.64
Storage						
WKSET	223.34	208.73	197.14	228.00	203.38	193.81
PGBLPGS	3565.00	3565.00	3565.00	3553.00	3553.00	3553.00
SHRPGS	326.16	324.21	357.74	408.05	417.58	438.90
Paging						
PAGE RATE	62.64	93.47	134.87	69.24	99.33	137.21
PAGE/CMD	11.96	16.04	20.91	13.13	16.74	21.48
I/0						
VIO RATE	43.18	48.31	53.39	41.46	47.33	50.46
VIO/CMD	8.24	8.29	8.28	7.86	7.98	7.90
MDSK/CMD (R)	2.793	2.715	2.713	2.468	2.444	2.432
Queueing						
PCTCPUQ	1.27	5.19	5.39	3.57	1.41	7.19
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	0.73	1.60	2.55	2.32	3.06	2.20
PCTIOQ	1.18	0.64	0.85	0.36	0.78	0.47
PRIVOPs						
PRIVOP RATE	189.13	214.01	240.48	102.48	117.25	127.31
PRIVOP/CMD	36.10	36.73	37.29	19.43	19.76	19.93
SFS Server						
WKSET	1553	1608	1611	1642	1637	1630
PAGE/CMD	0.00	0.00	0.14	0.00	0.15	0.13
PBT/CMD	0.014	0.016	0.021	0.019	0.020	0.021
CP/CMD	0.005	0.007	0.009	0.006	0.007	0.007
VIRT/CMD	0.008	0.009	0.012	0.013	0.014	0.014
FP REQ/CMD (Q)	2.90	2.94	2.93	2.28	2.23	2.33
IO/CMD (Q)	3.40	3.51	3.56	3.46	3.52	3.66
IO TIME/CMD (Q)	0.08	0.11	0.13	0.09	0.10	0.14
SFS TIME/CMD (Q)	0.13	0.16	0.20	0.13	0.15	0.19

Note: R=RESPONS2, Q=QUERY FILEPOOL STATUS, Unmarked=VMMAP, * = not a true ITR due to system constraints.

Table 20. 4381-13 / 16M / MAX% SFS

4.2.2.5 9370-80 / 16M / 35% SFS

1) WORKLOAD: FS7B35

2) HARDWARE CONFIGURATION

- PROCESSOR: 9370-80
- STORAGE: 16M

- DASD:

	PACK NAME	TYPE	CHANNEL
- SYSTEM:	VMSRES	9335	0C
	VMPK01	9335	0D
	VMPK02	9335	0E

	NUMBER OF PACKS	TYPE	NUMBER OF CHANNELS
- USER/PAGE:	4	9335	4

3) SOFTWARE CONFIGURATION

- DRIVER: FSID
- USER THINK TIME: RANDOM (Avg 10 Secs)
- CMS BLOCKSIZE: 1K
- USER VM SIZE: 2M
- USER PRIORITY: 64

- SERVER MACHINES:

	TYPE	NUMBER	VM SIZE	PRIORITY	OTHER OPTIONS
- SP 6:	SFS	1	16M	1	QDROP OFF USERS
- ESA 1.0:	SFS	1	16M	1	QDROP OFF USERS
	CRR	1	16M	64	

4) MEASUREMENT SUMMARY

This comparison includes a combination of Shared File System and EDF File System activity. The data shows equivalent storage requirements, a reduction in processor usage, and a reduction in virtual file I/Os.

This comparison demonstrates the same characteristics as those already described for the corresponding measurements made on the 4381 processor.

Average response time (AVG TOT (R)) for this environment changed by +4% to -5%.

	VM/SP 6			VM/ESA 1.0		
Real Storage						
TOTAL STORAGE	15916K			15916K		
NUCLEUS	429K			454K		
DYNAMIC	13864K			13816K		
FREE	1394K			1393K		
TRACE	228K			252K		
Run-ID	EC7601	EC7601	EC7601	EC7033	EC7033	EC7033
ACTIVE USERS	27	31	35	27	31	35
Response Time						
TRIV RESP	0.13	0.16	0.21	0.11	0.13	0.15
1ST RESP (R)	0.61	0.81	1.15	0.62	0.76	1.04
AVG TOT (R)	0.92	1.17	1.67	0.96	1.17	1.59
Throughput						
ITR	3.47	3.40	3.32	3.57	3.50	3.37
ETR (R)	2.433	2.754	2.982	2.472	2.755	2.992
CPU Usage						
PBT/CMD	0.288	0.294	0.301	0.280	0.286	0.297
CP/CMD	0.108	0.114	0.120	0.110	0.116	0.125
VIRT/CMD	0.180	0.180	0.181	0.171	0.170	0.172
PBT/CMD (R)	0.232	0.236	0.243	0.224	0.228	0.234
CP/CMD (R)	0.070	0.074	0.079	0.072	0.076	0.081
VIRT/CMD (R)	0.162	0.162	0.164	0.152	0.152	0.153
CPU Utilization						
TOTCPU	70.14	80.92	89.82	69.26	78.80	88.84
CPCPU	26.34	31.25	35.84	27.06	31.96	37.26
VIRTCPU	43.80	49.67	53.98	42.20	46.83	51.58
Storage						
WKSET	399.10	353.91	321.28	388.38	348.75	313.84
PGBLPGS	3458.00	3458.00	3458.00	3446.00	3446.00	3446.00
SHRPGS	351.57	366.02	364.24	438.37	441.06	448.61
Paging						
PAGE RATE	10.03	14.68	21.22	11.96	17.71	23.81
PAGE/CMD	4.12	5.33	7.12	4.84	6.43	7.96
I/0						
VIO RATE	22.24	25.03	27.12	21.70	24.10	26.31
VIO/CMD	9.14	9.09	9.09	8.78	8.75	8.79
MDSK/CMD (R)	3.611	3.571	3.584	3.309	3.302	3.293
Queueing						
PCTCPUQ	3.00	6.34	8.55	5.13	5.00	7.26
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	0.80	1.02	1.14	1.43	1.06	0.83
PCTIOQ	1.54	0.90	1.25	1.14	1.63	1.22
PRIVOPs						
PRIVOP RATE	108.72	122.84	132.55	66.57	75.01	81.92
PRIVOP/CMD	44.69	44.60	44.45	26.93	27.23	27.38
SFS Server						
WKSET	1364	1312	1305	1289	1277	1266
PAGE/CMD	0.00	0.00	0.00	0.00	0.00	0.00
PBT/CMD	0.023	0.027	0.026	0.036	0.037	0.038
CP/CMD	0.010	0.014	0.015	0.019	0.019	0.020
VIRT/CMD	0.013	0.013	0.012	0.018	0.018	0.018
FP REQ/CMD (Q)	2.09	2.10	2.09	1.49	1.47	1.48
IO/CMD (Q)	2.46	2.55	2.65	2.46	2.49	2.57
IO TIME/CMD (Q)	0.08	0.09	0.11	0.09	0.09	0.11
SFS TIME/CMD (Q)	0.13	0.14	0.17	0.12	0.13	0.15
Note: R=RESPONS2, Q=QUERY FILEPOOL STATUS, Unmarked=VMMAP						

Table 21. 9370-80 / 16M / 35% SFS

4.2.2.6 9370-30 / 8M / 35% SFS

```
1) WORKLOAD:     FS7B35

2) HARDWARE CONFIGURATION

   - PROCESSOR:  9370-30
   - STORAGE:    8M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  VMSRES     9332     0C
                 VMPK01     9332     0D
                 VMPK02     9332     0D

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       4        9332            2


3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    1K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64

   - SERVER MACHINES:
                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
            - SP 6:    SFS          1      16M         1    QDROP OFF USERS
            - ESA 1.0: SFS          1      16M         1    QDROP OFF USERS
                       CRR          1      16M        64

4) MEASUREMENT SUMMARY
```

This comparison includes a combination of Shared File System and EDF File System activity. The data shows similar storage requirements, a reduction in CPU usage, and a reduction in virtual file I/Os.

This environment demonstrates similar characteristics as the previous one. The reduction in both storage size and processing power did not reveal any unique performance trends.

Average response time (AVG TOT (R)) changed by -3% to -8%.

	VM/SP 6			VM/ESA 1.0		
Real Storage						
TOTAL STORAGE	8192K			8192K		
NUCLEUS	455K			479K		
DYNAMIC	6976K			6944K		
FREE	640K			640K		
TRACE	120K			128K		
Run-ID	E05437	E05438	E05440	EC7606	EC7606	EC7606
ACTIVE USERS	17	19	21	17	19	21
Response Time						
TRIV RESP	0.24	0.27	0.34	0.22	0.27	0.33
1ST RESP (R)	1.23	1.37	1.70	1.03	1.23	1.51
AVG TOT (R)	1.72	1.91	2.37	1.58	1.86	2.24
Throughput						
ITR	2.01	1.97	1.93	2.11	2.05	2.01
ETR (R)	1.438	1.577	1.695	1.466	1.591	1.709
CPU Usage						
PBT/CMD	0.496	0.506	0.519	0.474	0.487	0.499
CP/CMD	0.168	0.178	0.188	0.167	0.177	0.190
VIRT/CMD	0.328	0.329	0.331	0.307	0.310	0.309
PBT/CMD (R)	0.408	0.417	0.424	0.383	0.394	0.399
CP/CMD (R)	0.108	0.115	0.122	0.106	0.114	0.121
VIRT/CMD (R)	0.300	0.302	0.302	0.277	0.280	0.278
CPU Utilization						
TOTCPU	71.38	79.86	87.94	69.46	77.50	85.19
CPCPU	24.15	28.02	31.90	24.48	28.23	32.38
VIRTCPU	47.23	51.84	56.04	44.98	49.27	52.81
Storage						
WKSET	294.14	275.29	259.40	283.93	262.91	241.69
PGBLPGS	1737.00	1737.00	1737.00	1729.00	1729.00	1729.00
SHRPGS	294.24	294.94	290.87	361.18	363.75	364.37
Paging						
PAGE RATE	16.58	22.41	28.84	18.36	23.86	30.91
PAGE/CMD	11.53	14.21	17.01	12.52	15.00	18.09
I/0						
VIO RATE	13.12	14.37	15.49	12.88	13.97	14.96
VIO/CMD	9.12	9.11	9.14	8.79	8.78	8.75
MDSK/CMD (R)	3.598	3.636	3.580	3.293	3.344	3.313
Queueing						
PCTCPUQ	6.92	7.04	8.29	5.15	8.36	8.34
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	2.40	2.28	3.34	2.34	3.41	3.21
PCTIOQ	1.41	0.95	1.04	1.86	1.12	0.79
PRIVOPs						
PRIVOP RATE	39.88	42.77	45.79	16.97	18.31	19.71
PRIVOP/CMD	27.73	27.12	27.01	11.58	11.51	11.53
SFS Server						
WKSET	908	900	896	917	914	889
PAGE/CMD	0.15	0.13	0.25	0.14	0.27	0.37
PBT/CMD	0.045	0.047	0.050	0.053	0.053	0.054
CP/CMD	0.022	0.023	0.025	0.024	0.024	0.025
VIRT/CMD	0.023	0.024	0.025	0.028	0.028	0.029
FP REQ/CMD (Q)	2.11	2.10	2.09	1.47	1.49	1.47
IO/CMD (Q)	2.55	2.60	2.71	2.50	2.55	2.57
IO TIME/CMD (Q)	0.12	0.13	0.15	0.12	0.13	0.14
SFS TIME/CMD (Q)	0.18	0.19	0.21	0.17	0.18	0.20
Note: R=RESPONS2, Q=QUERY FILEPOOL STATUS, Unmarked=VMMAP						

Table 22. 9370-30 / 8M / 35% SFS

4.2.2.7 9370-30 / 4M / 0% SFS

```
1) WORKLOAD:     FS7B0

2) HARDWARE CONFIGURATION

   - PROCESSOR:  9370-30
   - STORAGE:    4M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  VMSRES     9332     0C
                 VMPK01     9332     0D
                 VMPK02     9332     0D

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       4        9332            2


3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    1K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64


4) MEASUREMENT SUMMARY
```

This comparison does not include any Shared File System activity, only the EDF File System was used. The data shows a reduction in storage requirements, a large reduction in CPU usage, and a slight reduction in virtual file I/Os.

The resulting improvements in this comparison are similar to those described in the 4381 EDF only measurements that have been previously discussed. However, this storage constrained environment takes more advantage of the storage improvements made for this release.

The net of this is seen by the large improvement in average response time as shown in the table below.

	VM/SP 6	VM/ESA 1.0
Real Storage		
TOTAL STORAGE	4096K	4096K
NUCLEUS	439K	463K
DYNAMIC	3260K	3232K
FREE	336K	336K
TRACE	60K	64K
Run-ID	EC6255	EC7602
ACTIVE USERS	17	17
Response Time		
TRIV RESP	0.73	0.53
1ST RESP (R)	1.85	1.41
AVG TOT (R)	2.45	1.92
Throughput		
ITR	* 1.89	* 2.10
ETR (R)	1.341	1.322
CPU Usage		
PBT/CMD	0.528	0.476
CP/CMD	0.222	0.197
VIRT/CMD	0.306	0.279
PBT/CMD (R)	0.459	0.414
CP/CMD (R)	0.157	0.138
VIRT/CMD (R)	0.302	0.276
CPU Utilization		
TOTCPU	70.84	62.86
CPCPU	29.83	26.02
VIRTCPU	41.01	36.83
Storage		
WKSET	176.65	173.68
PGBLPGS	809.00	802.00
SHRPGS	169.84	186.00
Paging		
PAGE RATE	53.96	41.59
PAGE/CMD	40.24	31.46
I/0		
VIO RATE	14.66	14.03
VIO/CMD	10.93	10.61
MDSK/CMD (R)	5.364	5.117
Queueing		
PCTCPUQ	5.71	4.29
PCTSTGQ	0.25	0.16
PCTPAGQ	9.11	7.27
PCTIOQ	2.41	1.52
PRIVOPs		
PRIVOP RATE	36.42	14.71
PRIVOP/CMD	27.16	11.13
Note: R=RESPONS2, Unmarked=VMMAP, * = not a true ITR due to system constraints		

Table 23. 9370-30 / 4M / 0% SFS

4.2.2.8 9370-30 / 768K / 0% SFS

```
1) WORKLOAD:     FS7B0

2) HARDWARE CONFIGURATION

   - PROCESSOR:  9370-30
   - STORAGE:    768K

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  VMSRES     9332     0C
                 VMPK01     9332     0D
                 VMPK02     9332     0D

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       4        9332            2


3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    1K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64


4) MEASUREMENT SUMMARY
```

The following results indicate that VM/ESA 1.0 370 feature can be run with CMS interactive users in 768K. The nucleus was generated with the small CP option. The small TOTCPU and high PAGE/CMD indicate a severely storage constrained environment. It should be noted that due to the storage constraints, these results could vary significantly with repeated measurements.

	VM/ESA 1.0
Real Storage	
TOTAL STORAGE	768K
NUCLEUS	296K
DYNAMIC	412K
FREE	47K
TRACE	12K
Run-ID	EC7034
ACTIVE USERS	3
Response Time	
TRIV RESP	1.52
1ST RESP (R)	5.86
AVG TOT (R)	8.24
Throughput	
ITR	* 0.91
ETR (R)	0.158
CPU Usage	
PBT/CMD	1.101
CP/CMD	0.809
VIRT/CMD	0.292
PBT/CMD (R)	0.788
CP/CMD (R)	0.502
VIRT/CMD (R)	0.286
CPU Utilization	
TOTCPU	17.39
CPCPU	12.78
VIRTCPU	4.62
Storage	
WKSET	271.52
PGBLPGS	84.00
SHRPGS	0.12
Paging	
PAGE RATE	37.18
PAGE/CMD	235.32
I/0	
VIO RATE	1.68
VIO/CMD	10.63
MDSK/CMD (R)	5.056
Queueing	
PCTCPUQ	0.56
PCTSTGQ	20.17
PCTPAGQ	23.81
PCTIOQ	1.12
PRIVOPs	
PRIVOP RATE	1.74
PRIVOP/CMD	11.01
Note: R=RESPONS2, Unmarked=VMMAP, * = not a true ITR due to system constraints.	

Table 24. 9370-30 / 768K / 0% SFS

4.3 EDF and SFS Comparisons

4.3.1 Introduction

The Regression section highlights what performance differences occur when going from the previous release to ESA 1.0 370 Feature. The measurements in this section compare the performance of EDF and SFS, when both are running on ESA 1.0 370 Feature.

SFS counts and timings obtained from the QUERY FILEPOOL STATUS command are provided in “Appendix B. SFS Counter Data” on page 155 to supplement the SFS measurement data provided in this section.

4.3.2 Measurements

4.3.2.1 4381-13 / 16M / EDF and SFS at Equal Utilization

```
1) WORKLOAD:     FS7B0, FS7B35, and FS7BMAX


2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380     01
                 PRF05      3380     01
                 PRF01      3380     02
                 PRF06      3380     02
                 PRFPRO     3380     03

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      10        3380            4


3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64

   - SERVER MACHINES:
                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
            - SP 6:    SFS          1      16M         1    QDROP OFF USERS
            - ESA 1.0: SFS          1      16M         1    QDROP OFF USERS
                       CRR          1      16M        64

4) MEASUREMENT SUMMARY
```

These measurements vary the SFS usage for three CPU utilizations: 65%, 75%, and 85%. For example, at 65% CPU utilization, there is a run for 0% SFS and a run for 35% SFS, both running on ESA 1.0 370 Feature. Note that the actual CPU utilization for the runs is close but does not exactly match the approximate values of 65%, 75%, and 85%. The columns of the table are arranged by CPU utilization to facilitate comparisons.

The number of users (ACTIVE USERS) decreased by 12% to 18% as the degree of SFS usage increased, for each of the fixed utilizations when compared to the 0% SFS runs. The processor busy time per command (PBT/CMD under CPU Usage) increased by 16% to 21% as the degree of SFS usage increased. Virtual file I/Os per command (FILE IO/CMD (Q)) also increased by 15% to 25%. These resource usage increases resulted in a response time (AVG TOT (R)) which increased an average of 0.1 seconds. This response time increase is minimal because the levels of resource contention are similar within each utilization. See section " 4381-13 / 16M / EDF and SFS with Equal Users" on page 104 to see what happens to response times for equal users (higher contention).

The storage requirements and response times increased between the 0% and 35% measurements but the increase did not continue as might be expected in going to the MAX% measurements (all files which SFS supports are in SFS). To achieve the MAX% SFS run while holding the CPU utilization at approximately 75%, there were four fewer users than for the 35% SFS run at this utilization. The working sets of these four users represented 23% of the pagable pages (PGBLPGS). Without these users, the paging per command (PAGE/CMD under Paging) dropped by 5 between the 35% and MAX% runs at 75% CPU utilization. This drop in paging accounts for the reduced response time between the 35% and MAX% measurements.

CPU Util. (Approx)	65%		75%			85%	
% SFS	0	35	0	35	MAX	0	35
Release	VM/ESA 1.0		VM/ESA 1.0			VM/ESA 1.0	
Real Storage							
TOTAL STORAGE	16384K		16384K			16384K	
NUCLEUS	488K		488K			488K	
DYNAMIC	14244K		14244K			14244K	
FREE	1395K		1395K			1395K	
TRACE	256K		256K			256K	
Run-ID	EC7603	EC7600	EC7603	EC7600	EC7031	EC7603	EC7600
ACTIVE USERS	69	59	77	67	63	85	75
Response Time							
TRIV RESP	0.23	0.28	0.29	0.36	0.29	0.38	0.47
1ST RESP (R)	0.42	0.49	0.55	0.65	0.56	0.70	0.83
AVG TOT (R)	0.62	0.70	0.80	0.89	0.81	0.98	1.14
Throughput							
ITR	9.77	8.18	9.41	8.06	7.75	* 9.05	* 7.74
ETR (R)	6.431	5.468	7.121	6.133	5.794	7.643	6.619
CPU Usage							
PBT/CMD	0.102	0.122	0.106	0.124	0.129	0.111	0.129
CP/CMD	0.026	0.034	0.029	0.036	0.036	0.031	0.039
VIRT/CMD	0.076	0.089	0.078	0.088	0.093	0.079	0.090
PBT/CMD (R)	0.097	0.101	0.099	0.102	0.101	0.102	0.105
CP/CMD (R)	0.021	0.022	0.022	0.023	0.023	0.023	0.025
VIRT/CMD (R)	0.076	0.079	0.077	0.079	0.078	0.079	0.080
CPU Utilization							
TOTCPU	65.85	66.85	75.73	76.11	74.76	84.44	85.55
CPCPU	16.82	18.40	20.28	21.98	21.07	23.73	25.89
VIRTCPU	49.03	48.45	55.45	54.13	53.69	60.71	59.66
Storage							
WKSET	204.18	220.98	189.50	207.66	212.14	178.96	192.83
PGBLPGS	3553.00	3553.00	3553.00	3553.00	3553.00	3553.00	3553.00
SHRPGS	312.21	398.79	327.42	404.37	408.53	324.89	411.84
Paging							
PAGE RATE	95.96	90.39	129.57	127.37	91.75	167.85	167.47
PAGE/CMD	14.92	16.53	18.20	20.77	15.84	21.96	25.30
I/0							
VIO RATE	65.42	46.75	72.62	52.09	45.62	78.40	56.32
VIO/CMD	10.17	8.55	10.20	8.49	7.87	10.26	8.51
MDSK/CMD (R)	4.700	3.048	4.676	3.005	2.425	4.712	3.015
FILE IO/CMD (Q)	4.700	5.478	4.676	5.395	5.825	4.712	5.425
Queueing							
PCTCPUQ	2.73	2.41	2.52	2.36	3.24	2.96	3.21
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	2.05	1.55	2.11	2.89	1.54	2.84	3.96
PCTIOQ	1.67	0.86	1.36	0.53	0.57	1.85	0.89
PRIVOPs							
PRIVOP RATE	68.39	90.53	74.83	98.98	114.18	79.67	105.80
PRIVOP/CMD	10.63	16.56	10.51	16.14	19.71	10.42	15.98
SFS Server							
WKSET	na	1458	na	1490	1643	na	1473
PAGE/CMD	na	0.00	na	0.14	0.15	na	0.27
PBT/CMD	na	0.014	na	0.013	0.020	na	0.014
CP/CMD	na	0.004	na	0.004	0.007	na	0.005
VIRT/CMD	na	0.009	na	0.009	0.014	na	0.009
FP REQ/CMD (Q)	na	1.52	na	1.49	2.25	na	1.47
IO/CMD (Q)	na	2.43	na	2.39	3.40	na	2.41
IO TIME/CMD (Q)	na	0.07	na	0.08	0.10	na	0.10
SFS TIME/CMD (Q)	na	0.10	na	0.12	0.15	na	0.14

Note: R=RESPONS2, Q=QUERY FILEPOOL STATUS, Unmarked=VMMAP, * = Not a true ITR due to system constraints.

Table 25. 4381-13 / 16M / EDF and SFS at Equal Utilization

4.3.2.2 4381-13 / 16M / EDF and SFS with Equal Users

1) WORKLOAD: FS7B0 and FS7B35

2) HARDWARE CONFIGURATION

- PROCESSOR: 4381-13
- STORAGE: 16M

- DASD:

	PACK NAME	TYPE	CHANNEL
- SYSTEM:	PRFRES	3380	01
	PRF05	3380	01
	PRF01	3380	02
	PRF06	3380	02
	PRFPRO	3380	03

	NUMBER OF PACKS	TYPE	NUMBER OF CHANNELS
- USER/PAGE:	10	3380	4

3) SOFTWARE CONFIGURATION

- DRIVER: FSID
- USER THINK TIME: RANDOM (Avg 10 Secs)
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2M
- USER PRIORITY: 64

- SERVER MACHINES:

	TYPE	NUMBER	VM SIZE	PRIORITY	OTHER OPTIONS
- SP 6:	SFS	1	16M	1	QDROP OFF USERS
- ESA 1.0:	SFS	1	16M	1	QDROP OFF USERS
	CRR	1	16M	64	

4) MEASUREMENT SUMMARY

Storage, processor, and response times all showed significant increases when holding the number of users constant.

The storage requirements (PAGE/CMD under Paging) showed an increase of 44% to 60% when going from the 0% SFS runs to the 35% SFS runs with an equal number of users. The increased paging is due to:

- commands which use SFS require/reference more pages and
- when the SFS server is running, as it would be in the 35% case, there is less real storage for user virtual machines.

The increased PAGE/CMD is a direct reflection of SFS′s greater real storage requirements. The PAGE/CMD increase in the equal utilization comparison (see “ 4381-13 / 16M / EDF and SFS at Equal Utilization” on page 101) was smaller because SFS was running fewer users.

SFS paging is low because the server is running with QDROP OFF and because, being a server, it frequently references its working set.

The processor busy time (PBT/CMD under CPU Usage) showed an increase of 24% for the 35% run. This increase was more than the measurements in section “ 4381-13 / 16M / EDF and SFS at Equal Utilization” on page 101 primarily because of the higher paging.

Average response time (AVG TOT (R)) increased from the corresponding EDF base measurements. The degree of increase depended on how constrained the system was to begin with. For the least constrained environment, the increase was 0.3 seconds. For the most constrained environment, the increase was 1.1 seconds. Nearly all the increase is due to higher levels of contention for the processor and real storage. This is shown by the equal utilization results, which indicate that

when the levels of contention are roughly equivalent, the SFS response time impact is about 0.1 seconds.

	VM/ESA 1.0 0% SFS			VM/ESA 1.0 35% SFS		
Real Storage						
TOTAL STORAGE	16384K			16384K		
NUCLEUS	488K			488K		
DYNAMIC	14244K			14244K		
FREE	1395K			1395K		
TRACE	256K			256K		
Run-ID	EC7603	EC7603	EC7603	EC7029	EC7029	EC7029
ACTIVE USERS	69	77	85	69	77	85
Response Time						
TRIV RESP	0.23	0.29	0.38	0.37	0.52	0.76
1ST RESP (R)	0.42	0.55	0.70	0.65	0.92	1.50
AVG TOT (R)	0.62	0.80	0.98	0.90	1.27	2.05
Throughput						
ITR	9.77	9.41	* 9.05	7.93	7.64	* 7.23
ETR (R)	6.431	7.121	7.643	6.210	6.753	6.886
CPU Usage						
PBT/CMD	0.102	0.106	0.111	0.126	0.131	0.138
CP/CMD	0.026	0.029	0.031	0.037	0.040	0.046
VIRT/CMD	0.076	0.078	0.079	0.089	0.091	0.093
PBT/CMD (R)	0.097	0.099	0.102	0.104	0.106	0.109
CP/CMD (R)	0.021	0.022	0.023	0.024	0.025	0.027
VIRT/CMD (R)	0.076	0.077	0.079	0.080	0.081	0.082
CPU Utilization						
TOTCPU	65.85	75.73	84.44	78.28	88.37	95.33
CPCPU	16.82	20.28	23.73	22.84	27.04	31.30
VIRTCPU	49.03	55.45	60.71	55.44	61.33	64.03
Storage						
WKSET	204.18	189.50	178.96	202.03	188.44	175.36
PGBLPGS	3553.00	3553.00	3553.00	3553.00	3553.00	3553.00
SHRPGS	312.21	327.42	324.89	405.21	423.63	423.11
Paging						
PAGE RATE	95.96	129.57	167.85	133.41	178.48	241.63
PAGE/CMD	14.92	18.20	21.96	21.48	26.43	35.09
I/0						
VIO RATE	65.42	72.62	78.40	52.90	57.58	58.74
VIO/CMD	10.17	10.20	10.26	8.52	8.53	8.53
MDSK/CMD (R)	4.700	4.676	4.712	3.021	3.016	2.962
Queueing						
PCTCPUQ	2.73	2.52	2.96	2.44	4.45	7.55
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	2.05	2.11	2.84	3.04	3.92	7.55
PCTIOQ	1.67	1.36	1.85	0.59	1.26	1.03
PRIVOPs						
PRIVOP RATE	68.39	74.83	79.67	101.16	109.28	110.48
PRIVOP/CMD	10.63	10.51	10.42	16.29	16.18	16.04
SFS Server						
WKSET	na	na	na	1451	1481	1536
PAGE/CMD	na	na	na	0.14	0.26	0.64
PBT/CMD	na	na	na	0.014	0.014	0.015
CP/CMD	na	na	na	0.005	0.005	0.005
VIRT/CMD	na	na	na	0.009	0.009	0.010
FP REQ/CMD(Q)	na	na	na	1.49	1.47	1.49
IO/CMD (Q)	na	na	na	2.44	2.49	2.58
IO TIME/CMD (Q)	na	na	na	0.08	0.11	0.16
SFS TIME/CMD (Q)	na	na	na	0.12	0.15	0.24

Note: R=RESPONS2, Q=QUERY FILEPOOL STATUS, Unmarked=VMMAP, * = Not a true ITR due to system constraints.

Table 26. 4381-13 / 16M / EDF and SFS with Equal Users

4.4 Connectivity

4.4.1 Introduction

This section describes measurements comparing ESA 1.0 370 Feature to SP 6 for the three types of APPC connectivity (local APPC/VM, AVS, and TSAF). The purpose of making these measurements was to insure that performance of APPC communications had not degraded significantly from the previous release. There is a slight degradation in most of these measurements. This degradation has two sources:

1. Small increase (about 5%) in CP IUCV pathlength due to CRR;

2. Small increase in CP pathlength due to >16M support.

Some of the TSAF environments show improvement despite the extra CP pathlength.

The Connectivity Driver tool does little besides APPCVM communications, so the runs described in this section are very sensitive to increases in IUCV pathlength.

VMMAP, RESPONS2, FSID, and the Connectivity Driver tool were used to produce this data. In all the tables in this section, items of the form ″xxx/CMD″ (for example, ″PBT/CMD″) have a slightly different meaning than similar variables elsewhere in this report. In all cases, ″CMD″ here means not one user interaction, but one APPCVM SEND/RECEIVE pair. A single user interaction will, in general, cause several of these SEND/RECEIVE pairs to occur.

4.4.2 APPC/VM Measurements

4.4.2.1 Multi-user 512 byte

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-92E  CPU 0 only
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      15        3380            4

   - TAPE:       MONITOR  3480 on channel 9

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - REQUEST SIZE:     512 bytes
   - TRANSACTION:      11 requests of 512 bytes each
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   RES1         1       3M         1    QDROP OFF USERS
                                                            FAVORED, FAVORED 100
```

4) MEASUREMENT SUMMARY

The following table shows the results of the 512-byte multi-user local APPC runs. The CPU time for a single APPC 512-byte transaction (PBT/CMD) has increased by about 7.5% in the new release because of the increased pathlength in IUCV and extended storage support mentioned above. In a non-paging environment extended storage doesn't affect performance, while the added pathlength for this support has performance costs.

	VM/SP 6	VM/ESA 1.0
Real Storage		
TOTAL STORAGE	16384K	16384K
NUCLEUS	464K	492K
DYNAMIC	14292K	14244K
FREE	236K	256K
TRACE	1392K	1392K
Run-ID	1	2
ACTIVE USERS	150	150
Response Time		
1ST RESP (R)	0.053	0.067
Throughput		
ETR	29.660	29.360
CPU Usage		
PBT/CMD	0.00196	0.00212
CP/CMD	0.00174	0.00187
VIRT/CMD	0.00022	0.00025
CPU Utilization		
TOTCPU	64.06	68.55
CPCPU	56.88	60.53
VIRTCPU	7.18	8.02
Paging		
PAGE RATE	0.00	0.00
PAGE/CMD	0.000	0.000
I/O		
VIO RATE	0.00	0.00
VIO/CMD	0.0000	0.0000
Queuing		
PCTCPUQ	0.24	0.28
PCTSTGQ	0.00	0.00
PCTPAGQ	0.00	0.00
PCTIOQ	0.00	0.00
PRIVOPs		
PRIVOP RATE	1303.57	1290.79
PRIVOP/CMD	3.995	3.997
Note: R=RESPONS2, unmarked=VMMAP		

Table 27. Local Multi-user: APPC/VM 4381-92E (one processor only) **512-byte**

4.4.2.2 Multi-user 4096 byte

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-92E  CPU 0 only
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
        USER/PAGE:      15        3380            4

   - TAPE:       MONITOR  3480 on channel 9

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - REQUEST SIZE:     4096 bytes
   - TRANSACTION:      10 requests of 4096 bytes each
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   RES1         1       3M         1    QDROP OFF USERS
                                                            FAVORED, FAVORED 100

4) MEASUREMENT SUMMARY
```

The following table shows the results of the 4096-byte multi-user local APPC runs. The CPU time for a single APPC 4096-byte transaction (PBT/CMD) has increased by about 6.5% in the new release for the reasons mentioned in the “Introduction” on page 107. The degradation is slightly less in this 4096-byte case because the increased IUCV pathlength due to CRR is amortized over a larger data transfer size.

	VM/SP 6	VM/ESA 1.0
Real Storage		
TOTAL STORAGE	16384K	16384K
NUCLEUS	464K	492K
DYNAMIC	14292K	14244K
FREE	236K	256K
TRACE	1392K	1392K
Run-ID	3	4
ACTIVE USERS	150	150
Response Time		
1ST RESP (R)	0.054	0.065
Throughput		
ETR	29.730	29.402
CPU Usage		
PBT/CMD	0.00217	0.00232
CP/CMD	0.00193	0.00206
VIRT/CMD	0.00023	0.00026
CPU Utilization		
TOTCPU	64.40	68.30
CPCPU	57.43	60.59
VIRTCPU	6.97	7.70
Paging		
PAGE RATE	0.00	0.00
PAGE/CMD	0.000	0.000
I/O		
VIO RATE	0.00	0.00
VIO/CMD	0.0000	0.0000
Queuing		
PCTCPUQ	0.21	0.20
PCTSTGQ	0.00	0.00
PCTPAGQ	0.00	0.00
PCTIOQ	0.00	0.00
PRIVOPs		
PRIVOP RATE	1189.16	1175.96
PRIVOP/CMD	4.000	4.000
Note: R=RESPONS2, unmarked=VMMAP		

Table 28. Local Multi-user: **APPC/VM** 4381-92E (one processor only) **4096-byte**

4.4.2.3 Multi-user 102400 byte

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-92E  CPU 0 only
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      15        3380            4

   - TAPE:       MONITOR  3480 on channel 9

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - REQUEST SIZE:     102400 bytes
   - TRANSACTION:      5 requests of 102400 bytes each
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   RES1         1       3M         1    QDROP OFF USERS
                                                            FAVORED, FAVORED 100

4) MEASUREMENT SUMMARY
```

The following table shows the results of the 102400-byte multi-user local APPC runs. The CPU time for a single APPC 102400-byte transaction (PBT/CMD) has increased by about 4.9% in the new release for the reasons mentioned in “Introduction” on page 107. The additional CP pathlength due to CRR plays a still smaller role here, when compared to the smaller transaction sizes. However, pathlength increase from the extended storage support plays a significantly greater role in the 102400-byte case because of the greater number of user pages involved in each transaction.

	VM/SP 6	VM/ESA 1.0
Real Storage		
TOTAL STORAGE	16384K	16384K
NUCLEUS	464K	492K
DYNAMIC	14292K	14244K
FREE	236K	256K
TRACE	1392K	1392K
Run-ID	5	6
ACTIVE USERS	100	100
Response Time		
1ST RESP (R)	0.137	0.157
Throughput		
ETR	19.113	19.040
CPU Usage		
PBT/CMD	0.00701	0.00737
CP/CMD	0.00668	0.00700
VIRT/CMD	0.00033	0.00037
CPU Utilization		
TOTCPU	66.98	70.14
CPCPU	63.86	66.63
VIRTCPU	3.12	3.51
Paging		
PAGE RATE	9.20	11.96
PAGE/CMD	0.096	0.126
I/O		
VIO RATE	0.00	0.00
VIO/CMD	0.0000	0.0000
Queuing		
PCTCPUQ	0.26	0.23
PCTSTGQ	0.00	0.00
PCTPAGQ	0.05	0.07
PCTIOQ	0.00	0.00
PRIVOPs		
PRIVOP RATE	382.29	380.84
PRIVOP/CMD	4.000	4.000
Note: R=RESPONS2, unmarked=VMMAP		

Table 29. Local Multi-user: APPC/VM 102400-byte 4381-92E (one processor only)

4.4.3 AVS Measurements

4.4.3.1 Multi-user 512 byte

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   USER SIDE (APPLICATION/REQUESTOR)

   - PROCESSOR:  4381-92E  CPU 0 only
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      15        3380            4

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

   RESOURCE SIDE (SERVER)

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  VMSRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       2        3380            1

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - REQUEST SIZE:     512 bytes
   - TRANSACTION:      11 requests of 512 bytes each
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   VTAM/AVS     1       8M         1    QDROP OFF USERS
                                                            FAVORED, FAVORED 100
                                                            DIAGNOSE X'98'

4) MEASUREMENT SUMMARY
```

The following table shows the results of the 512-byte multi-user runs through AVS. The degradation from SP 6 is less than 3% in processor busy time per transaction. The increase in processor time is mostly in CP, again mostly due to the small increase in APPCVM pathlength.

	VM/SP 6		VM/ESA 1.0	
Side	User	Resource	User	Resource
Real Storage				
TOTAL STORAGE	16384K	16384K	16384K	16384K
NUCLEUS	464K	468K	492K	492K
DYNAMIC	14292K	14288K	14244K	14244K
FREE	236K	236K	256K	256K
TRACE	1392K	1392K	1392K	1392K
Run-ID	7	7	8	8
ACTIVE USERS	36	na	36	na
Response Time				
1ST RESP (R)	1.001	na	1.031	na
Throughput				
ETR	5.928	na	5.858	na
CPU Usage				
PBT/CMD	0.0097	0.0132	0.0099	0.0136
CP/CMD	0.0025	0.0034	0.0026	0.0038
VIRT/CMD	0.0073	0.0097	0.0073	0.0098
CPU Utilization				
TOTCPU	63.58	85.91	63.97	87.89
CPCPU	16.04	22.38	16.85	24.62
VIRTCPU	47.53	63.52	47.12	63.27
Paging				
PAGE RATE	0.00	0.00	0.00	0.00
PAGE/CMD	0.000	0.000	0.000	0.000
I/O				
VIO RATE	58.29	77.32	57.04	76.54
VIO/CMD	0.8939	1.1857	0.8851	1.1877
Queuing				
PCTCPUQ	0.41	4.96	0.50	5.37
PCTSTGQ	0.00	0.00	0.00	0.00
PCTPAGQ	0.00	0.00	0.00	0.00
PCTIOQ	0.00	0.00	0.00	0.00
PRIVOPs				
PRIVOP RATE	1.10	588.45	1.10	582.57
PRIVOP/CMD	0.017	9.024	0.017	9.040
VTAM/AVS				
WKSET	na	na	na	na
TOT CPU/CMD	0.0087	0.0118	0.0088	0.0121
CP CPU/CMD	0.0015	0.0022	0.0016	0.0025
VIRT CPU/CMD	0.0072	0.0096	0.0072	0.0096
NON SP/CMD	0.89	1.21	0.88	1.21
PAGE RD/SEC	0.00	0.00	0.00	0.00
Note: R=RESPONS2, unmarked=VMMAP				

Table 30. VTAM/AVS 512-byte Multi-user: 4381-92E (one processor only) to 4381-13

4.4.3.2 Multi-user 4096 byte

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   USER SIDE (APPLICATION/REQUESTOR)

   - PROCESSOR:  4381-92E  CPU 0 only
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      15        3380            4

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

   RESOURCE SIDE (SERVER)

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  VMSRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       2        3380            1

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - REQUEST SIZE:     4096 bytes
   - TRANSACTION:      10 requests of 4096 bytes each
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   VTAM/AVS     1       8M         1    QDROP OFF USERS
                                                            FAVORED, FAVORED 100
                                                            DIAGNOSE X'98'

4) MEASUREMENT SUMMARY
```

The following table shows the results of the 4096-byte multi-user runs through AVS. The degradation from SP 6 is less than 2% in processor busy time per transaction, for the reasons noted previously in the “Introduction” on page 107.

	VM/SP 6		VM/ESA 1.0	
Side	User	Resource	User	Resource
Real Storage				
TOTAL STORAGE	16384K	16384K	16384K	16384K
NUCLEUS	464K	468K	492K	492K
DYNAMIC	14292K	14288K	14244K	14244K
FREE	236K	236K	256K	256K
TRACE	1392K	1392K	1392K	1392K
Run-ID	9	9	10	10
ACTIVE USERS	25	na	25	na
Response Time				
1ST RESP (R)	0.953	na	1.024	na
Throughput				
ETR	4.048	na	4.093	na
CPU Usage				
PBT/CMD	0.0149	0.0166	0.0153	0.0168
CP/CMD	0.0041	0.0045	0.0042	0.0048
VIRT/CMD	0.0108	0.0122	0.0111	0.0120
CPU Utilization				
TOTCPU	60.32	67.36	62.64	68.78
CPCPU	16.47	18.13	17.09	19.55
VIRTCPU	43.86	49.24	45.55	49.23
Paging				
PAGE RATE	0.00	0.00	0.00	0.00
PAGE/CMD	0.000	0.000	0.000	0.000
I/O				
VIO RATE	56.37	64.15	54.47	61.72
VIO/CMD	1.3924	1.5846	1.3307	1.5078
Queuing				
PCTCPUQ	0.06	5.33	0.13	4.96
PCTSTGQ	0.00	0.00	0.00	0.00
PCTPAGQ	0.00	0.00	0.00	0.00
PCTIOQ	0.00	0.00	0.00	0.00
PRIVOPs				
PRIVOP RATE	1.11	371.93	1.07	369.77
PRIVOP/CMD	0.027	9.187	0.026	9.033
VTAM/AVS				
WKSET	na	na	na	na
TOT CPU/CMD	0.0135	0.0153	0.0138	0.0151
CP CPU/CMD	0.0027	0.0031	0.0028	0.0032
VIRT CPU/CMD	0.0107	0.0122	0.0110	0.0119
NON SP/CMD	1.39	1.64	1.33	1.53
PAGE RD/SEC	0.00	0.00	0.00	0.00
Note: R=RESPONS2, unmarked=VMMAP				

Table 31. VTAM/AVS 4096-byte Multi-user: 4381-92E (one processor only) to 4381-13

4.4.3.3 Multi-user 102400 byte

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   USER SIDE (APPLICATION/REQUESTOR)

   - PROCESSOR:  4381-92E  CPU 0 only
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      15        3380            4

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

   RESOURCE SIDE (SERVER)

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  VMSRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       2        3380            1

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - REQUEST SIZE:     102400 bytes
   - TRANSACTION:      7 requests of 102400 bytes each
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   VTAM/AVS     1       8M         1    QDROP OFF USERS
                                                            FAVORED, FAVORED 100
                                                            DIAGNOSE X'98'

4) MEASUREMENT SUMMARY
```

The following table shows the results of the 102400-byte multi-user runs through AVS. In this case, The degradation from SP 6 is a little less than 6% in processor busy time per transaction. Again, this increase is largely due to pathlength increase in CP due to the extended storage support.

	VM/SP 6		VM/ESA 1.0	
Side	User	Resource	User	Resource
Real Storage				
TOTAL STORAGE	16384K	16384K	16384K	16384K
NUCLEUS	464K	468K	492K	492K
DYNAMIC	14292K	14288K	14244K	14244K
FREE	236K	236K	256K	256K
TRACE	1392K	1392K	1392K	1392K
Run-ID	11	11	12	12
ACTIVE USERS	6	na	6	na
Response Time				
1ST RESP (R)	4.576	na	5.305	na
Throughput				
ETR	0.607	na	0.582	na
CPU Usage				
PBT/CMD	0.1319	0.0910	0.1399	0.0966
CP/CMD	0.0293	0.0240	0.0304	0.0265
VIRT/CMD	0.1026	0.0671	0.1095	0.0701
CPU Utilization				
TOTCPU	56.02	38.66	56.97	39.34
CPCPU	12.43	10.18	12.38	10.79
VIRTCPU	43.59	28.48	44.59	28.55
Paging				
PAGE RATE	0.00	0.00	0.00	0.00
PAGE/CMD	0.000	0.000	0.000	0.000
I/O				
VIO RATE	12.30	9.65	11.72	9.05
VIO/CMD	2.8964	2.2724	2.8784	2.2227
Queuing				
PCTCPUQ	0.24	0.00	0.48	0.41
PCTSTGQ	0.00	0.00	0.00	0.00
PCTPAGQ	0.00	0.00	0.00	0.00
PCTIOQ	0.00	0.00	0.00	0.00
PRIVOPs				
PRIVOP RATE	1.13	64.96	1.11	62.84
PRIVOP/CMD	0.266	15.297	0.273	15.433
VTAM/AVS				
WKSET	na	na	na	na
TOT CPU/CMD	0.1256	0.0871	0.1330	0.0916
CP CPU/CMD	0.0232	0.0194	0.0237	0.0206
VIRT CPU/CMD	0.1024	0.0677	0.1093	0.0710
NON SP/CMD	2.90	2.30	2.87	2.27
PAGE RD/SEC	0.00	0.00	0.00	0.00
Note: R=RESPONS2, unmarked=VMMAP				

Table 32. VTAM/AVS 102400-byte Multi-user: 4381-92E (one processor only) to 4381-13

4.4.4 TSAF Measurements

The following TSAF measurements required the application of a fix to TSAF that will be shipped as an APAR. The number of this APAR is not available at this writing.

The small transaction sizes (512 and 4096 bytes) show a 4-6% increase in CP time per transaction. The virtual time in the 512-byte case has improved enough to show a 4-6% improvement in total processor time per command. The 4096-byte case shows this improvement only on the resource side. The large transaction case (102400 bytes) shows improvement even in CP time, due largely to a decrease in the number of I/O's per command.

4.4.4.1 Multi-user 512 byte

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   USER SIDE (APPLICATION/REQUESTOR)

   - PROCESSOR:  4381-92E  CPU 0 only
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      15        3380            4

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

   RESOURCE SIDE (SERVER)

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  VMSRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       2        3380            1

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - REQUEST SIZE:     512 bytes
   - TRANSACTION:      11 requests of 512 bytes each
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   TSAF         1       4M         1    QDROP OFF USERS
                                                            FAVORED, FAVORED 100

4) MEASUREMENT SUMMARY
```

The following table shows the results of the 512-byte multi-user runs using TSAF. Although the CP time per transaction is increased because of the slightly longer CP APPC pathlength, this increase is masked by a large reduction in TSAF virtual time. The end result is a 4-6% reduction in total CPU per transaction.

	SP 6		ESA 1.0 370 Feature	
Side	User	Resource	User	Resource
Real Storage				
TOTAL STORAGE	16384K	16384K	16384K	16384K
NUCLEUS	464K	468K	492K	492K
DYNAMIC	14292K	14288K	14244K	14244K
FREE	236K	236K	256K	256K
TRACE	1392K	1392K	1392K	1392K
Run-ID	13	13	14	14
ACTIVE USERS	30	na	30	na
Response Time				
1ST RESP (R)	0.223	na	0.222	na
Throughput				
ETR	5.560	na	5.642	na
CPU Usage				
PBT/CMD	0.0069	0.0082	0.0065	0.0079
CP/CMD	0.0039	0.0045	0.0041	0.0047
VIRT/CMD	0.0030	0.0037	0.0024	0.0031
CPU Utilization				
TOTCPU	42.29	49.88	40.32	48.76
CPCPU	23.98	27.40	25.26	29.39
VIRTCPU	18.31	22.47	15.07	19.38
Paging				
PAGE RATE	0.00	0.00	0.70	0.00
PAGE/CMD	0.000	0.000	0.011	0.000
I/O				
VIO RATE	0.99	0.91	0.95	1.68
VIO/CMD	0.0162	0.0149	0.0153	0.0271
Queuing				
PCTCPUQ	0.38	3.72	0.27	6.61
PCTSTGQ	0.00	0.00	0.00	0.00
PCTPAGQ	0.00	0.00	0.00	0.00
PCTIOQ	0.00	0.00	0.05	0.00
PRIVOPs				
PRIVOP RATE	92.97	502.40	92.36	514.58
PRIVOP/CMD	1.520	8.215	1.488	8.292
TSAF MACHINE				
WKSET	na	na	na	na
TOT CPU/CMD	0.0056	0.0064	0.0051	0.0059
CP CPU/CMD	0.0026	0.0030	0.0027	0.0031
VIRT CPU/CMD	0.0029	0.0034	0.0024	0.0028
NON SP/CMD	2.05	2.07	2.00	2.05
PAGE RD/SEC	0.00	0.00	0.00	0.00
Note: R=RESPONS2, unmarked=VMMAP				

Table 33. TSAF 512-byte multi-user: 4381-92E (one processor only) to 4381-13

4.4.4.2 Multi-user 4096 byte

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   USER SIDE (APPLICATION/REQUESTOR)

   - PROCESSOR:  4381-92E  CPU 0 only
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      15        3380            4

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

   RESOURCE SIDE (SERVER)

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  VMSRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       2        3380            1

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - REQUEST SIZE:     4096 bytes
   - TRANSACTION:      10 requests of 4096 bytes each
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   TSAF         1       4M         1    QDROP OFF USERS
                                                            FAVORED, FAVORED 100

4) MEASUREMENT SUMMARY
```

The following table shows the results of the 4096-byte multi-user runs using TSAF. The CP time per command is up because of lengthened CP APPC pathlength, but on the resource side the TSAF virtual pathlength reduction more than makes up for the increase. The TSAF virtual CPU for the user side is up slightly.

	SP 6		ESA 1.0 370 Feature	
Side	User	Resource	User	Resource
Real Storage				
TOTAL STORAGE	16384K	16384K	16384K	16384K
NUCLEUS	464K	468K	492K	492K
DYNAMIC	14292K	14288K	14244K	14244K
FREE	236K	236K	256K	256K
TRACE	1392K	1392K	1392K	1392K
Run-ID	17	17	18	18
ACTIVE USERS	25	na	25	na
Response Time				
1ST RESP (R)	0.269	na	0.278	na
Throughput				
ETR	4.705	na	4.780	na
CPU Usage				
PBT/CMD	0.0080	0.0099	0.0084	0.0096
CP/CMD	0.0046	0.0052	0.0048	0.0055
VIRT/CMD	0.0034	0.0047	0.0036	0.0041
CPU Utilization				
TOTCPU	37.86	46.78	39.95	45.86
CPCPU	21.75	24.57	22.93	26.44
VIRTCPU	16.10	22.22	17.02	19.41
Paging				
PAGE RATE	0.00	0.00	0.55	0.00
PAGE/CMD	0.000	0.000	0.012	0.000
I/O				
VIO RATE	0.32	5.91	0.25	4.07
VIO/CMD	0.0068	0.1256	0.0052	0.0851
Queuing				
PCTCPUQ	0.38	4.13	0.06	3.72
PCTSTGQ	0.00	0.00	0.00	0.00
PCTPAGQ	0.00	0.00	0.00	0.00
PCTIOQ	0.00	0.41	0.00	0.00
PRIVOPs				
PRIVOP RATE	96.08	387.66	93.77	398.02
PRIVOP/CMD	2.042	8.239	1.962	8.327
TSAF MACHINE				
WKSET	na	na	na	na
TOT CPU/CMD	0.0065	0.0081	0.0070	0.0075
CP CPU/CMD	0.0032	0.0036	0.0034	0.0038
VIRT CPU/CMD	0.0034	0.0045	0.0036	0.0038
NON SP/CMD	2.03	2.22	2.06	2.20
PAGE RD/SEC	0.00	0.00	0.00	0.00
Note: R=RESPONS2, unmarked=VMMAP				

Table 34. TSAF 4096-byte multi-user: 4381-92E (one processor only) to 4381-13

4.4.4.3 Multi-user 102400 byte

```
1) WORKLOAD:     Multi-user

2) HARDWARE CONFIGURATION

   USER SIDE (APPLICATION/REQUESTOR)

   - PROCESSOR:  4381-92E  CPU 0 only
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      15        3380            4

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

   RESOURCE SIDE (SERVER)

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  VMSRES     3380      1

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:       2        3380            1

   - TAPE:       MONITOR  3480 on channel 9

   - CTCA:       3088 adapter on channel 5

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 5 Secs)
   - REQUEST SIZE:     102400 bytes
   - TRANSACTION:      7 requests of 102400 bytes each
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64 (Users)
   - USER PRIORITY:    1 (Resources)

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   TSAF         1       4M         1    QDROP OFF USERS
                                                            FAVORED, FAVORED 100

4) MEASUREMENT SUMMARY
```

The following table shows the results of the 102400-byte multi-user runs using TSAF. The virtual CPU per command on the resource side has been reduced by 25%, and the virtual CPU per command on the user side is the same as in VM/SP 6. Even with the increased CP pathlength, CP CPU is down because of the substantially smaller number of I/O's per command. Overall there is a 2% reduction in CPU per command on the user side and an 18% reduction on the resource side.

	SP 6		ESA 1.0 370 Feature	
Side	User	Resource	User	Resource
Real Storage				
TOTAL STORAGE	16384K	16384K	16384K	16384K
NUCLEUS	464K	468K	492K	492K
DYNAMIC	14292K	14288K	14244K	14244K
FREE	236K	236K	256K	256K
TRACE	1392K	1392K	1392K	1392K
Run-ID	15	15	16	16
ACTIVE USERS	10	na	10	na
Response Time				
1ST RESP (R)	2.940	na	2.404	na
Throughput				
ETR	1.263	na	1.320	na
CPU Usage				
PBT/CMD	0.0686	0.0848	0.0672	0.0691
CP/CMD	0.0340	0.0421	0.0326	0.0373
VIRT/CMD	0.0346	0.0427	0.0346	0.0319
CPU Utilization				
TOTCPU	43.32	53.56	44.35	45.62
CPCPU	21.45	26.58	21.51	24.59
VIRTCPU	21.87	26.98	22.84	21.03
Paging				
PAGE RATE	0.00	0.00	0.27	0.00
PAGE/CMD	0.000	0.000	0.041	0.000
I/O				
VIO RATE	1.15	6.21	0.49	10.43
VIO/CMD	0.1821	0.9831	0.0742	1.5803
Queuing				
PCTCPUQ	0.30	0.41	0.76	0.00
PCTSTGQ	0.00	0.00	0.00	0.00
PCTPAGQ	0.00	0.00	0.00	0.00
PCTIOQ	0.00	0.41	0.00	0.00
PRIVOPs				
PRIVOP RATE	90.35	146.23	94.63	122.64
PRIVOP/CMD	14.303	23.150	14.338	18.582
TSAF MACHINE				
WKSET	na	na	na	na
TOT CPU/CMD	0.0604	0.0766	0.0588	0.0614
CP CPU/CMD	0.0259	0.0335	0.0242	0.0293
VIRT CPU/CMD	0.0346	0.0431	0.0346	0.0321
NON SP/CMD	16.60	18.12	11.09	13.50
PAGE RD/SEC	0.00	0.00	0.00	0.00
Note: R=RESPONS2, unmarked=VMMAP				

Table 35. TSAF 102400-byte multi-user: 4381-92E (one processor only) to 4381-13

4.5 Coordinated Resource Recovery

4.5.1 Introduction

The measurements and analysis for CRR were made in order to quantify CRR performance and determine limitations. The workloads used were designed to stress CRR functions and avoid non-CRR function overhead where possible. CRR function requests are made at a far higher rate in these measurements than is expected of non-test systems.

System paging requirements were kept to a minimum in the measurements by design. Therefore, storage related performance measures, such as page rate, are not as meaningful in these runs. Some of these variables have been marked as 'na' in the tables that follow.

All measurements shown are of multi-user environments. The workloads involved are made up of an application that exercises CRR function. This application is called the CRRTOOL and is described in more detail in "CRRTOOL" on page 178. The difference in some workloads is from variations in the parameters used for the CRRTOOL. The following workloads are used:

- FSCRR1 - Two filepools are updated and work committed. Each command consists of two iterations of this process.
- FSCRR3 - Two filepools are updated and work committed. Each command consists of four iterations of this process.
- FSCRR4 - One filepool is updated and a protected conversation used, then work is committed. Each command consists of two iterations.
- FSCRR5 - One filepool is updated and another is read; work is committed. Each command consists of two iterations.

The following measurements all include data on the CRR Recovery Server, which can be found in the tables under 'CRR Machine'. In addition, there is data specific to CRR which was obtained through use of the QUERY FILEPOOL STATUS command. Included are:

- *SYNCPT RATE* - syncpoint rate is the number of times work was committed per second through the two-phase commit process. This will differ from ETR (external throughput rate) in some runs because the "commands" counted for external throughput are the invocations of the CRRTOOL and some invocations result in multiple commits.
- *LOG REQ RATE* - CRR log write request rate is the number of log write requests made to the CRR Recovery Server per second.
- *LOG I/O RATE* - CRR log I/O rate is the rate per second of physical I/Os for logging. CRR uses dual logging (i.e. logs to identical disks) and this rate includes both disks.
- *CHECKPT RATE* - CRR checkpoint rate is the count of CRR checkpoints performed per second. Checkpoint processing is used in the log disk wrapping maintenance.
- *LOG I/O TIME* - CRR log I/O time is the average elapsed time per log I/O processing in milliseconds.

4.5.2 Measurements

4.5.2.1 4381-13 / 16M / CPU Utilization Runs

```
1) WORKLOAD:     FSCRR1

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380     01

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      13        3380            4

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   CRR          1      16M        1     QDROP OFF USERS
                       SFS          3      16M        1     QDROP OFF USERS

4) MEASUREMENT SUMMARY
```

This set of measurements shows the performance of CRR at various levels of CPU utilization. The different CPU utilizations were obtained by varying the number of users. The workload used here results in the users performing syncpoints which result in two-phase commits.

The performance benefit realized from optimizations in the Recovery Server are illustrated here. Relative performance improves at higher levels of utilization. Key points that reflect this are :

- The PBT/CMD decreases as the number of users increases, particularly in the CRR Recovery Server. From the 6 user run to the 60 user run, PBT/CMD in the Recovery Server decreases over 10%. The majority of the improvement is in CP CPU time.
- The Log I/Os per command decrease as the number of users increases. They decreased from 8 per command (or 4 per syncpoint) to 5.4 per command (2.7 per syncpoint).

	ESA 1.0 370 Feature					
Real Storage						
TOTAL STORAGE			16384K			
NUCLEUS			488K			
DYNAMIC			14244K			
FREE			1395K			
TRACE			256K			
Run-ID	11	12	13	14	15	16
ACTIVE USERS	6	36	42	48	54	60
Response Time						
TRIV RESP	0.00	0.00	0.00	0.00	0.00	0.00
1ST RESP (R)	na	na	na	na	na	na
AVG TOT (R)	0.293	0.509	0.614	0.743	0.984	1.312
Throughput						
ITR	5.656	5.750	5.794	5.770	5.821	5.858
ETR (R)	0.583	3.420	3.969	4.462	4.943	5.326
CPU Usage						
PBT/CMD	0.1768	0.1739	0.1726	0.1733	0.1718	0.1707
CP/CMD	0.0784	0.0742	0.0726	0.0727	0.0711	0.0698
VIRT/CMD	0.0986	0.0997	0.0999	0.1005	0.1007	0.1008
PBT/CMD (R)	0.048	0.049	0.049	0.049	0.049	0.049
CP/CMD (R)	0.020	0.020	0.020	0.020	0.020	0.020
VIRT/CMD (R)	0.028	0.029	0.029	0.029	0.029	0.029
CPU Utilization						
TOTCPU	10.31	59.48	68.50	77.33	84.94	90.91
CPCPU	4.57	25.38	28.83	32.46	35.15	37.20
VIRTCPU	5.75	34.10	39.67	44.86	49.78	53.71
Storage						
WKSET	1201.47	332.03	245.87	233.90	195.11	212.82
PGBLPGS	3553.00	3553.00	3553.00	3553.00	3553.00	3553.00
SHRPGS	298.00	259.01	205.76	212.93	216.00	231.24
Paging						
PAGE RATE	0.00	0.12	0.03	0.21	0.05	0.25
PAGE/CMD	0.000	0.035	0.008	0.047	0.010	0.047
I/O						
VIO RATE	0.00	0.00	0.00	0.00	0.00	0.01
VIO/CMD	0.000	0.000	0.000	0.000	0.000	0.002
MDSK/CMD (R)	0.000	0.000	0.000	0.000	0.000	0.00
Queuing						
PCTCPUQ	0.66	2.64	2.83	3.09	2.96	3.61
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	0.00	0.00	0.00	0.00	0.00	0.000
PCTIOQ	0.00	0.00	0.00	0.00	0.00	0.000
PRIVOPs						
PRIVOP RATE	41.55	237.39	271.91	307.81	336.45	355.83
PRIVOP/CMD	71.269	69.412	68.508	68.985	68.066	66.810
CRR MACHINE						
WKSET	2352	781	867	856	818	969
TOT CPU/CMD	0.0286	0.0276	0.0272	0.0266	0.0259	0.0255
CP/CMD	0.0152	0.0143	0.0140	0.0137	0.0130	0.0127
VIRT/CMD	0.0133	0.0133	0.0132	0.0129	0.0128	0.0127
NON SP/CMD	8.0	7.1	6.7	6.5	6.0	5.4
PAGE RD/SEC	0.0	0.0	0.0	0.0	0.0	0.0
Coordination (Q)						
SYNCPT RATE	1.160	6.843	7.956	8.941	9.915	10.681
LOG REQ RATE	3.479	20.528	23.877	26.826	29.742	32.053
LOG I/O RATE	4.639	24.749	26.941	29.328	29.861	29.007
CHECKPT RATE	0.000	0.000	0.000	0.000	0.004	0.002
LOG I/O TIME	7.460	9.223	9.773	10.183	11.216	12.451
Note: R=RESPONS2, Unmarked=VMMAP, Q=QUERY FILEPOOL STATUS						

Table 36. 4381-13 / 16M / CPU Utilization Runs: Number of users varied to create different CPU utilizations.

4.5.2.2 4381-13 / 16M / Concurrency Analysis Runs

```
1) WORKLOAD:     FSCRR1 and FSCRR3

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380     01

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      13        3380            4

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs for runs 14, 15, and 16
                               Avg  5 Secs for runs 31 and 32.)
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   CRR          1      16M        1     QDROP OFF USERS
                       SFS          3      16M        1     QDROP OFF USERS

4) MEASUREMENT SUMMARY
```

These measurements illustrate the effect of concurrency of syncpoint processing on system performance. In measurements 31 and 32, the syncpoint rates were obtained by changing the number of users, their think time, and the iteration value for CRRTOOLU. The iteration value determines the number of times per command that work is committed. In runs 14, 15, and 16, there are two syncpoints per command while there are four syncpoints per command for runs 31 and 32. This methodology provides data for runs with similar syncpoint rates, but different levels of concurrency. It is important to note that the users in these workloads have small working sets, which minimizes the performance impact due to storage constraints. In other environments, performance could degrade with more users because of greater system paging demands.

The results show that the higher the degree of concurrency, the higher the benefit from Recovery Server optimizations. The degree of concurrency is lower in runs 31 and 32 than runs 14, 15, and 16. As seen in earlier measurements, the relative performance improves as the syncpoint rate increases. Expected values for CRR CPU time and I/O rate are extrapolated from the higher concurrency runs for syncpoint rates achieved in the lower concurrency runs. The extrapolations are made so that performance can be compared at equal syncpoint rates. By comparing the expected values with the actual values for the lower concurrency runs, one can see there is a slight performance degradation associated with the different concurrency level.

	CRR CPU / Syncpoint		Log I/Os / Syncpoint	
Run	Expected	Actual	Expected	Actual
31	13.1 msec	13.3 msec	3.11	3.30
32	12.7 msec	12.8 msec	2.70	3.02

	ESA 1.0 370 Feature				
Real Storage					
TOTAL STORAGE			16384K		
NUCLEUS			488K		
DYNAMIC			14244K		
FREE			1395K		
TRACE			256K		
Run-ID	14	31	15	16	32
ACTIVE USERS	48	15	54	60	18
Think Time	10	5	10	10	5
Iterations	2	4	2	2	4
Response Time					
TRIV RESP	0.00	0.00	0.00	0.00	0.00
1ST RESP (R)	na	na	na	na	na
AVG TOT (R)	0.743	1.276	0.984	1.312	1.589
Throughput					
ITR	5.770	3.010	5.819	5.858	3.052
ETR (R)	4.462	2.383	4.943	5.326	2.682
CPU Usage					
PBT/CMD	0.1733	0.3322	0.1718	0.1707	0.3276
CP/CMD	0.0727	0.1367	0.0711	0.0698	0.1314
VIRT/CMD	0.1005	0.1955	0.1007	0.1008	0.1963
PBT/CMD (R)	0.049	0.092	0.049	0.049	0.091
CP/CMD (R)	0.020	0.035	0.020	0.020	0.034
VIRT/CMD (R)	0.029	0.057	0.029	0.029	0.057
CPU Utilization					
TOTCPU	77.33	79.16	84.94	90.91	87.87
CPCPU	32.46	32.58	35.15	37.20	35.23
VIRTCPU	44.86	46.58	49.78	53.71	52.64
Storage					
WKSET	233.90	593.67	195.11	212.82	431.88
PGBLPGS	3553.00	3553.00	3553.00	3553.00	3553.00
SHRPGS	212.93	336.00	216.00	231.24	206.00
Paging					
PAGE RATE	0.21	0.00	0.05	0.25	0.01
PAGE/CMD	0.047	0.000	0.010	0.047	0.004
I/O					
VIO RATE	0.00	0.00	0.00	0.01	0.00
VIO/CMD	0.000	0.000	0.000	0.002	0.000
MDSK/CMD (R)	0.000	0.000	0.000	0.00	0.000
Queuing					
PCTCPUQ	3.09	6.98	2.96	3.61	6.23
PCTSTGQ	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	0.00	0.00	0.00	0.000	0.00
PCTIOQ	0.00	0.00	0.00	0.000	0.00
PRIVOPs					
PRIVOP RATE	307.81	322.58	336.45	355.83	354.49
PRIVOP/CMD	68.985	135.367	68.066	66.810	132.174
CRR MACHINE					
WKSET	856	2975	818	969	2849
TOT CPU/CMD	0.0266	0.0532	0.0259	0.0255	0.0510
CP/CMD	0.0137	0.0275	0.0130	0.0127	0.0257
VIRT/CMD	0.0129	0.0256	0.0128	0.0127	0.0253
NON SP/CMD	6.5	13.1	6.0	5.4	12.0
PAGE RD/SEC	0.00	0.00	0.00	0.0	0.00
Coordination (Q)					
SYNCPT RATE	8.941	9.562	9.915	10.681	10.726
LOG REQ RATE	26.826	28.689	29.742	32.053	32.181
LOG I/O RATE	29.328	31.527	29.861	29.007	32.363
CHECKPT RATE	0.000	0.001	0.004	0.002	0.003
LOG I/O TIME	10.183	9.972	11.216	12.451	10.656
Note: R=RESPONS2, Unmarked=VMMAP, Q=QUERY FILEPOOL STATUS					

Table 37. 4381-13 / 16M / CRR Concurrency Analysis Runs: The degree of concurrency in a system affects CRR performance.

4.5.2.3 4381-13 / 16M / Various CRR Environment Runs

```
1) WORKLOAD:     FSCRR1, FSCRR4, and FSCRR5

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-13
   - STORAGE:    16M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380     01

                 NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      13        3380            4

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64

                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
   - SERVER MACHINE:   CRR          1      16M        1     QDROP OFF USERS
                       SFS          3      16M        1     QDROP OFF USERS
                       RESOURCE     1       3M        1     FAVORED, FAVORED 100
                                                            QDROP OFF USERS

4) MEASUREMENT SUMMARY
```

These measurements show the CRR performance in various environments. Three workloads were used:

1. FSCRR1 - Each user updates two filepools and then commits the changes. This results in a two-phase commit.
2. FSCRR4 - Each user updates one filepool and uses a protected conversation connected to a server machine running the CRRTOOLR program. Syncpoint processing here results in a two-phase commit.
3. FSCRR5 - Each user updates one filepool and reads from another filepool. The syncpoint that occurs here is a one-phase (or simple) commit.

It is important to note that the CRR Recovery Server is not involved in one-phase commit processing, as it is in the two-phase commit. For more information on the differences, see “Coordinated Resource Recovery” on page 32. Compared to the two-phase commit, the one-phase commit requires significantly less resources. PBT/CMD is less than 50% of the two-phase commit. Time spent in CP and the privop rate are significantly less.

The throughput and PBT/CMD of the two-phase commit runs (16 and 41) are similar. In run 41, the users switched from updating one of the filepools to using a protected conversation with the CRRTOOLR server machine. Each syncpoint in the user machine results in another syncpoint in the CRRTOOLR server machine. Because of this, the protected conversation run results in over twice the overhead in the CRR Recovery Server as the run without a protected conversation. The system privop rate increases as well. Since the overhead of the CRRTOOLR server machine is less than that of the SFS server it replaces, the system overheads of the two runs are nearly equivalent.

	ESA 1.0 370 Feature		
Real Storage			
TOTAL STORAGE		16384K	
NUCLEUS		488K	
DYNAMIC		14244K	
FREE		1395K	
TRACE		256K	
Run-ID	42	16	41
ACTIVE USERS	60	60	60
Environment	One-phase	Two-phase	Two-phase PCA
Response Time			
TRIV RESP	0.00	0.00	0.00
1ST RESP (R)	na	na	na
AVG TOT (R)	0.185	1.312	1.192
Throughput			
ITR	13.162	5.858	5.880
ETR (R)	5.927	5.326	5.387
CPU Usage			
PBT/CMD	0.0760	0.1707	0.1701
CP/CMD	0.0292	0.0698	0.0818
VIRT/CMD	0.0468	0.1008	0.0882
PBT/CMD (R)	0.026	0.049	0.040
CP/CMD (R)	0.011	0.020	0.019
VIRT/CMD (R)	0.015	0.029	0.021
CPU Utilization			
TOTCPU	45.03	90.91	91.61
CPCPU	17.29	37.20	44.08
VIRTCPU	27.74	53.71	47.53
Storage			
WKSET	204.36	212.82	205.19
PGBLPGS	3553.00	3553.00	3553.00
SHRPGS	198.00	231.24	149.00
Paging			
PAGE RATE	0.23	0.25	0.32
PAGE/CMD	0.039	0.047	0.059
I/O			
VIO RATE	0.00	0.01	0.00
VIO/CMD	0.000	0.002	0.000
MDSK/CMD (R)	0.000	0.00	0.000
Queuing			
PCTCPUQ	0.88	3.61	3.80
PCTSTGQ	0.00	0.00	0.00
PCTPAGQ	0.00	0.000	0.00
PCTIOQ	0.00	0.000	0.00
PRIVOPs			
PRIVOP RATE	154.32	355.83	558.57
PRIVOP/CMD	26.037	66.810	103.689
CRR MACHINE			
WKSET	608	969	809
TOT CPU/CMD	0.0000	0.0255	0.0575
CP/CMD	0.0000	0.0127	0.0289
VIRT/CMD	0.0000	0.0127	0.0287
NON SP/CMD	0.0	5.4	14.0
PAGE RD/SEC	0.00	0.0	0.0
Coordination (Q)			
SYNCPT RATE	na	10.681	21.580
LOG REQ RATE	na	32.053	75.534
LOG I/O RATE	na	29.007	76.230
CHECKPT RATE	na	0.002	0.003
LOG I/O TIME	na	12.451	8.885
Note: R=RESPONS2, Unmarked=VMMAP, Q=QUERY FILEPOOL STATUS			

Table 38. 4381-13 / 16M / Various CRR Environment Runs: Measurements show overhead of coordination in different environments.

4.6 GCS

4.6.1 Introduction

This section describes measurements comparing ESA 1.0 370 feature running GCS6 to ESA 1.0 370 feature running GCS370. The purpose of making these measurements was simply to insure that performance of GCS had not degraded significantly from the previous release.

The results show GCS370 to be equivalent to GCS6.

4.6.2 Measurements

4.6.2.1 4381-13 / 16M / GCS6 vs. GCS370

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

- PROCESSOR: 4381-13
- STORAGE: 16M

- DASD:

	PACK NAME	TYPE	CHANNEL
- SYSTEM:	PRFRES	3380	01
	PRF01	3380	02
	PRFPRO	3380	03
	PRF05	3380	03
	PRF06	3380	03

	NUMBER OF PACKS	TYPE	NUMBER OF CHANNELS
- USER/PAGE	10	3380	4

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES	TYPE	LINESPEED
3705	1	16	SDLC	9600
3725	1	16	SDLC	9600

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- USER THINK TIME: RANDOM (Avg 10 Secs)
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2M
- USER PRIORITY: 64

	TYPE	NUMBER	VM SIZE	PRIORITY	OTHER OPTIONS
- SERVER MACHINE:	VTAM/VSCS	1	12M	1	QDROP, FAVOR
	SFS	1	16M	1	QDROP OFF USERS
	CRR	1	16M	64	

4) MEASUREMENT SUMMARY

The measurement data shows that GCS370 is equivalent to GCS6.

	ESA 1.0/GCS6			ESA 1.0/GCS370		
Real Storage						
TOTAL STORAGE	16384K			16384K		
NUCLEUS	488K			488K		
DYNAMIC	14244K			14244K		
FREE	1396K			1396K		
TRACE	256K			256K		
Run-ID	dr137	dr137	dr137	dr137	dr137	dr137
ACTIVE USERS	59	67	75	59	67	75
Response Time						
TRIV RESP	0.36	0.51	0.70	0.36	0.50	0.71
1ST RESP (T)	0.66	0.94	1.28	0.68	0.93	1.28
AVG TOT (T)	1.21	1.62	2.15	1.24	1.59	2.14
Throughput						
ITR	*7.19	*6.99	*6.85	*7.14	*7.04	*6.85
ETR (T)	5.27	5.80	6.13	5.24	5.81	6.16
CPU Usage						
PBT/CMD	0.139	0.143	0.146	0.140	0.142	0.146
CP/CMD	0.044	0.047	0.050	0.044	0.047	0.050
VIRT/CMD	0.095	0.096	0.096	0.095	0.095	0.096
PBT/CMD (R)	na	na	na	na	na	na
CP/CMD (R)	na	na	na	na	na	na
VIRT/CMD (R)	na	na	na	na	na	na
CPU Utilization						
TOTCPU	73.25	82.69	89.45	73.22	82.57	90.02
CPCPU	23.15	27.13	30.44	23.21	27.07	30.64
VIRTCPU	50.10	55.57	59.02	50.01	55.50	59.38
Storage						
WKSET	258.63	237.71	220.86	253.98	234.21	216.94
PGBLPGS	3469.00	3469.00	3469.00	3469.00	3469.00	3469.00
SHRPGS	259.04	363.31	372.95	379.11	378.57	388.13
Paging						
PAGE RATE	97.30	134.25	170.55	97.94	133.43	172.38
PAGE/CMD	18.46	23.15	27.82	18.69	22.97	27.98
I/O						
VIO RATE	64.77	69.55	72.00	64.50	69.48	72.41
VIO/CMD	12.29	12.09	11.75	12.31	11.96	11.75
MDSK/CMD (R)	na	na	na	na	na	na
Queuing						
PCTCPUQ	4.69	7.21	9.33	4.85	6.80	9.56
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	4.08	5.49	6.64	4.36	5.16	6.34
PCTIOQ	1.23	1.22	1.32	1.27	1.25	1.39
PRIVOPs						
PRIVOP RATE	138.53	154.16	162.92	138.52	153.60	164.30
PRIVOP/CMD	26.29	26.58	26.58	26.44	26.43	26.67
VTAM MACHINE						
WKSET	757	790	797	759	784	791
TOT CPU/CMD	0.019	0.018	0.017	0.019	0.018	0.017
CP/CMD	0.007	0.006	0.006	0.007	0.006	0.006
VIRT/CMD	0.012	0.012	0.011	0.012	0.012	0.011
NON SP/CMD	8.748	7.844	7.194	8.721	7.780	7.208
PAGE RD/SEC	na	na	na	na	na	na

Note: T=TPNS, R=RESPONS2, Unmarked=VMMAP, * = Not a true ITR due to system constraints.

Table 39. 4381-13 / 16 M / GCS6 to GCS370

4.7 Greater Than 16M of Real Storage

4.7.1 Introduction

All ESA 1.0 370 Feature measurements for extended storage support were conducted on a 4381-13 with 32M of real storage. These results were compared to SP 6 with 16M of real storage and to ESA 1.0 370 Feature with 16M of real storage. Three intervals were run at 70%, 80%, and 90% CPU utilization with the FS7B35 workload.

4.7.2 Measurements

4.7.2.1 4381-13 / 32M / 35% SFS

```
1) WORKLOAD:     FS7B35

2) HARDWARE CONFIGURATION

   - PROCESSOR:  4381-13
   - STORAGE:    32M

   - DASD:       PACK NAME  TYPE  CHANNEL
                 ---------  ----  -------
      - SYSTEM:  PRFRES     3380     01
                 PRF05      3380     01
                 PRF01      3380     02
                 PRF06      3380     02
                 PRFPRO     3380     03

      - USER:    NUMBER OF PACKS  TYPE  NUMBER OF CHANNELS
                 ---------------  ----  ------------------
      - USER/PAGE:      10        3380            4

3) SOFTWARE CONFIGURATION

   - DRIVER:           FSID
   - USER THINK TIME:  RANDOM (Avg 10 Secs)
   - CMS BLOCKSIZE:    4K
   - USER VM SIZE:     2M
   - USER PRIORITY:    64

   - SERVER MACHINES:
                       TYPE      NUMBER  VM SIZE  PRIORITY  OTHER OPTIONS
                       ----      ------  -------  --------  -------------
            - SP 6:    SFS          1      16M         1    QDROP OFF USERS
            - ESA 1.0: SFS          1      16M         1    QDROP OFF USERS
                       CRR          1      16M        64

4) MEASUREMENT SUMMARY
```

Following the measurement summary there are two tables. The first table shows a comparison of ESA 1.0 370 Feature with 16M to ESA 1.0 370 Feature with 32M. The second table shows a comparison of SP 6 with 16M to ESA 1.0 370 Feature with 32M.

As expected the greater than 16M support reduces the average response time because of a reduction in the number of page I/Os. Compared to ESA 1.0 370 Feature with 16M, the average response time decreased 34% to 46% depending on the level of system loading.

Overall processor busy time per command decreased. This is due to a large savings in page fault handling, which results in lower CP CPU consumption. There is a slight increase in supervisor CPU time for the movement of pages across the 16M boundary line and to check the location of pages. Compared to ESA 1.0 370 Feature with 16M, PBT/CMD decreased 8% to 12%. Also, CP/CMD decreased 24% to 28% while VIRT/CMD remained relatively constant.

Comparing ESA 1.0 370 Feature with 32M to SP 6 with 16M, the average response time decreased 39% to 53%. PBT/CMD decreased 13% to 17%, while CP/CMD decreased 21% to 28%. Also, VIRT/CMD decreased 11% to 14%.

Since the measurements below were compared to only 32M of real storage, some further improvements would be observed if 64M of real storage were used.

	VM/ESA 1.0			VM/ESA 1.0		
Real Storage						
TOTAL STORAGE	16384K			32680K		
NUCLEUS	488K			552K		
DYNAMIC	14244K			30412K		
FREE	1395K			1395K		
TRACE	256K			320K		
Run-ID	EC7600	EC7600	EC7600	EC7605	EC7605	EC7605
ACTIVE USERS	59	67	75	59	67	75
Response Time						
TRIV RESP	0.28	0.36	0.47	0.13	0.14	0.16
1ST RESP (R)	0.49	0.65	0.83	0.29	0.33	0.39
AVG TOT (R)	0.70	0.89	1.14	0.46	0.52	0.61
Throughput						
ITR	8.18	8.06	* 7.74	8.92	8.77	8.80
ETR (R)	5.468	6.133	6.619	5.598	6.275	6.943
CPU Usage						
PBT/CMD	0.122	0.124	0.129	0.112	0.114	0.114
CP/CMD	0.034	0.036	0.039	0.026	0.027	0.028
VIRT/CMD	0.089	0.088	0.090	0.086	0.087	0.086
PBT/CMD (R)	0.101	0.102	0.105	0.094	0.096	0.095
CP/CMD (R)	0.022	0.023	0.025	0.018	0.019	0.019
VIRT/CMD (R)	0.079	0.079	0.080	0.076	0.077	0.076
CPU Utilization						
TOTCPU	66.85	76.11	85.55	62.75	71.51	78.94
CPCPU	18.40	21.98	25.89	14.79	17.05	19.29
VIRTCPU	48.45	54.13	59.66	47.96	54.46	59.64
Storage						
WKSET	220.98	207.66	192.83	461.12	410.88	370.40
PGBLPGS	3553.00	3553.00	3553.00	3520.00	3520.00	3520.00
SHRPGS	398.79	404.37	411.84	336.16	346.26	371.42
Paging						
PAGE RATE	90.39	127.37	167.47	20.99	25.68	34.46
PAGE/CMD	16.53	20.77	25.30	3.75	4.09	4.96
I/0						
VIO RATE	46.75	52.09	56.32	47.64	53.36	59.07
VIO/CMD	8.55	8.49	8.51	8.51	8.50	8.51
MDSK/CMD (R)	3.048	3.005	3.015	3.025	3.067	2.994
Queueing						
PCTCPUQ	2.41	2.36	3.21	0.78	1.60	1.91
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	1.55	2.89	3.96	0.69	0.15	0.48
PCTIOQ	0.86	0.53	0.89	0.52	0.38	0.34
PRIVOPs						
PRIVOP RATE	90.53	98.98	105.80	92.71	105.96	114.39
PRIVOP/CMD	16.56	16.14	15.98	16.56	16.89	16.48
SFS Server						
WKSET	1458	1490	1473	3110	2539	2450
PAGE/CMD	0.00	0.14	0.27	0.00	0.00	0.00
PBT/CMD	0.014	0.013	0.014	0.013	0.014	0.013
CP/CMD	0.004	0.004	0.005	0.004	0.005	0.005
VIRT/CMD	0.009	0.009	0.009	0.009	0.009	0.009
FP REQ/CMD (Q)	1.52	1.49	1.47	1.45	1.48	1.47
IO/CMD (Q)	2.43	2.39	2.41	2.33	2.40	2.43
IO TIME (Q)	0.07	0.08	0.10	0.05	0.05	0.06
SFS TIME (Q)	0.10	0.12	0.14	0.08	0.08	0.09

Note: R=RESPONS2, Q=QUERY FILEPOOL STATUS, Unmarked=VMMAP, * = Not a true ITR due to system constraints.

Table 40. 4381-13 / 32M / 35% SFS (ESA 1.0 370 Feature 16M vs. 32M)

	VM/SP 6			VM/ESA 1.0		
Real Storage						
TOTAL STORAGE		16384K			32680K	
NUCLEUS		463K			552K	
DYNAMIC		14292K			30412K	
FREE		1392K			1395K	
TRACE		236K			320K	
Run-ID	EC4292	EC4292	EC4292	EC7605	EC7605	EC7605
ACTIVE USERS	59	67	75	59	67	75
Response Time						
TRIV RESP	0.31	0.40	0.57	0.13	0.14	0.16
1ST RESP (R)	0.53	0.72	0.99	0.29	0.33	0.39
AVG TOT (R)	0.75	0.97	1.31	0.46	0.52	0.61
Throughput						
ITR	7.81	7.57	* 7.31	8.92	8.77	8.80
ETR (R)	5.398	6.018	6.478	5.598	6.275	6.943
CPU Usage						
PBT/CMD	0.128	0.132	0.137	0.112	0.114	0.114
CP/CMD	0.033	0.036	0.039	0.026	0.027	0.028
VIRT/CMD	0.095	0.097	0.098	0.086	0.087	0.086
PBT/CMD (R)	0.107	0.111	0.114	0.094	0.096	0.095
CP/CMD (R)	0.021	0.024	0.025	0.018	0.019	0.019
VIRT/CMD (R)	0.086	0.087	0.089	0.076	0.077	0.076
CPU Utilization						
TOTCPU	69.10	79.48	88.60	62.75	71.51	78.94
CPCPU	17.59	21.34	25.23	14.79	17.05	19.29
VIRTCPU	51.51	58.14	63.37	47.96	54.46	59.64
Storage						
WKSET	228.43	210.21	196.88	461.12	410.88	370.40
PGBLPGS	3565.00	3565.00	3565.00	3520.00	3520.00	3520.00
SHRPGS	319.53	332.05	329.63	336.16	346.26	371.42
Paging						
PAGE RATE	87.17	125.24	169.97	20.99	25.68	34.46
PAGE/CMD	16.15	20.81	26.24	3.75	4.09	4.96
I/0						
VIO RATE	47.97	53.50	57.22	47.64	53.36	59.07
VIO/CMD	8.89	8.89	8.83	8.51	8.50	8.51
MDSK/CMD (R)	3.315	3.329	3.298	3.025	3.067	2.994
Queueing						
PCTCPUQ	3.16	3.79	4.00	0.78	1.60	1.91
PCTSTGQ	0.00	0.00	0.00	0.00	0.00	0.00
PCTPAGQ	2.11	1.78	3.31	0.69	0.15	0.48
PCTIOQ	0.88	0.85	0.69	0.52	0.38	0.34
PRIVOPs						
PRIVOP RATE	189.21	202.29	218.40	92.71	105.96	114.39
PRIVOP/CMD	35.05	33.61	33.71	16.56	16.89	16.48
SFS Server						
WKSET	1480	1501	1514	3110	2539	2450
PAGE/CMD	0.00	0.15	0.27	0.00	0.00	0.00
PBT/CMD	0.011	0.013	0.013	0.013	0.014	0.013
CP/CMD	0.005	0.006	0.005	0.004	0.005	0.005
VIRT/CMD	0.006	0.007	0.008	0.009	0.009	0.009
FP REQ/CMD (Q)	2.11	2.09	2.10	1.45	1.48	1.47
IO/CMD (Q)	2.36	2.43	2.50	2.33	2.40	2.43
IO TIME (Q)	0.07	0.08	0.11	0.05	0.05	0.06
SFS TIME (Q)	0.10	0.13	0.16	0.08	0.08	0.09

Note: R=RESPONS2, Q=QUERY FILEPOOL STATUS, Unmarked=VMMAP, * = Not a true ITR due to system constraints.

Table 41. 4381-13 / 32M / 35% SFS (SP 6 16M vs. ESA 1.0 370 Feature 32M)

Appendix A. CMS Trace Data

A.1.1 Measurement Methodology

The technique used to collect CMS trace data involves issuing selected commands from a 3270 terminal while the VM/370 Program Event Recording (PER) facility is active. Only the CMS (virtual) instruction execution path is traced. CP paths are not included.

Next, the TRREAD (Trace Read) program reads the trace file created by the VM/370 PER facility. It extracts the essential data and writes that data into the designated CMS output file. This file is then ready to be analyzed by the STARS program.

The STARS (System Trace Analysis Reports) program produces reports containing information about the instruction path taken while performing a given function. The primary input to STARS is the instruction trace CMS file (created by TRREAD). Another input is a storage map file which defines the virtual storage area used by CMS. This corresponds to the load map created during the CMS build process.

The main purpose of STARS is to relate the instruction execution shown by the trace to the virtual storage map and to break down the traced path into the separate contributions made by each CMS module. STARS produces a file containing a module scenario report, instruction distribution report, reference distribution report, a machine interrupt report, and an instruction mix report. All of this information is taken into account during the analysis of this trace data.

Traces were made with CMS running in a 370 virtual machine and also in an XA virtual machine. It is important to note that a 3277 terminal was used for the 370 mode traces, while a 3278 terminal was used for XA mode traces. The impact will be explained in further detail in the sections below.

A.1.2 Enhanced Disk Format (EDF) Commands Traced

A set of twenty-five EDF commands were traced for both SP6 and ESA1. These commands were chosen to exercise a large percentage of the common functions executed in CMS. The functions measured include EXEC processing, XEDIT related commands, program management, storage management, file system activity, and OS simulation. The data for these commands will show the effect of the new release on previously architected function (regression). The following are the EDF commands that were traced:

```
SET       BLIP OFF
ACCESS    294 B/A
COPY      MASTER SCRIPT A MASTER FILE A
COMPARE   MASTER SCRIPT A MASTER FILE A
RENAME    MASTER FILE A TEST SCRIPT A (UPDIRT
LISTFILE  TEST SCRIPT A (LABEL
XEDIT     TEST SCRIPT A
 XXXX      (invalid command)
 LOCATE    /EXECUTION
 CHANGE    /EXECUTION/DEFINITION
 NEXT      2
 INPUT
           This is an input test line.(eob)
           (eob)
 DELETE    1
 TOP
 FILE
ERASE     TEST SCRIPT A
EXEC2     X Y (ALL
QUERY     DISK
FILEDEF   IN DISK DUMMY FILE
PRINT     MASTER SCRIPT A
ASSEMBLE  BR14
LOAD      BR14
GENMOD    BR14
BR14
REXREX    X Y (ALL
```

A.1.3 Shared File System (SFS) Commands Traced

For SFS we trace CMS commands that would typically be issued by a user of SFS file pools. They are similar in function to the above EDF commands. In addition, several commands which are specific to a user of a file pool are included. File pool specific commands that we measure are CREATE ALIAS, GRANT AUTHORITY, and REVOKE AUTHORITY. The SFS commands traced are broken down by the virtual machine to which the overhead is charged. A typical SFS command will show virtual machine activity in both the user and server machines. The following are the SFS commands that were traced:

```
ACCESS   OPERATOR.F100  B/A
COPY     MASTER SCRIPT A MASTER FILE A
RENAME   MASTER FILE A TEST SCRIPT A (UPDIRT
XEDIT    TEST SCRIPT A
 FILE
ERASE    TEST SCRIPT A
ASSEMBLE BR14
LOAD     BR14
GENMOD   BR14
BR14
CREATE   ALIAS MASTER SCRIPT . PHANTOM = .FIRST
GRANT    AUTHORITY MASTER SCRIPT . TO OPERATOR
REVOKE   AUTHORITY MASTER SCRIPT . FROM  OPERATOR
XEDIT    BR14 LISTING A
 FILE
FILEDEF IN READER
```

Notes:

- These commands show up in our tables prefixed either with S to denote server virtual machine overhead or U to denote user virtual machine activity.
- The second XEDIT command is issued to show pathlength differences associated with the CRR machine. The first XEDIT command will show instruction overhead to initialize Co-ordinated Resource Recovery. This is a one-time occurrence which does not show up in sub-sequent XEDITs. The overhead is not directly associated with XEDIT. It occurs there due to the sequence of commands traced. The second XEDIT is referred to as XEDIT2 in our tabular data to distinguish it from its predecessor.

- The FILEDEF command is a non-SFS command in VM/ESA 1.0, while in VM/SP 6 there is an SFS interaction. The FILEDEF command was added for the analysis of CRR impact to EDF commands that are issued after SFS activity (which does interact with CRR). Analysis shows this impact to be negligible.

A.1.4 Pathlength Comparisons

This section contains the virtual pathlength data for the commands traced. The first set of trace data contains commands that are using the EDF File System. The second set of trace data contains the user side of commands that use the Shared File System. The third set of traces contains the SFS server side for the same commands. Also shown in the tables are the deltas for running an ESA 1.0 virtual machine in XA mode, as compared to running it in 370 mode on the 370 Feature.

A.1.4.1 EDF Pathlength Comparisons

The following table contains the pathlength data for the set of commands where the files being used are in the EDF File System. Included here is the virtual pathlength executed in the *user's* virtual machine.

The general trend shown in the pathlength table below indicates a small increase in pathlength for the average command. This increase was offset by the large reduction in pathlength for the COPY command, netting a slight improvement in the overall average for the twenty five commands.

The XA mode data shows an average increase of about 6% over running in 370 mode. This is primarily to support the different virtual machine structure, control blocks, and to provide the capability to use storage above 16 Megabytes. Part of this increase includes the additional overhead required to handle I/O interrupt processing. Since the majority of the commands have the same number of I/O interrupts, this processing is reflected as a flat increase to all commands. As seen in the table below, some commands increased by much more, or much less, than the average 6% increase. This is due to this flat increase which is a larger percentage hit for smaller commands. The two modes were traced using different terminal types. A 3278, which supports extended attributes, was used for the XA mode traces, while a 3277 was used for 370 mode. Additional processing is required by CMS for extended attributes when in fullscreen scenarios such as the XEDIT and XEDIT subcommands. This explains the larger increase seen in these commands.

Also traced, but not shown in the table, was a REXX EXEC which made heavy usage of EXECIO. Improvements were made in this area, and tracing was used to verify and quantify the benefit. The SP6 pathlength was 72,744 and the ESA 1.0 pathlength was 62,744. This EXEC improved by 14% due to reduced overhead in the linkage between REXX and EXECIO.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
SET	8,606	8,820	214	9,939	1,119
ACCESS	16,720	17,876	1,156	20,360	2,484
COPY	69,462	29,137	-40,325	30,283	1,146
COMPARE	59,089	66,954	7,865	70,086	3,132
RENAME	20,715	21,393	678	22,817	1,424
LISTFILE	10,236	10,400	164	12,842	2,442
XEDIT	65,747	73,693	7,946	78,644	4,951
XXXX	28,106	28,912	806	30,977	2,065
LOCATE	43,490	44,587	1,097	50,076	5,489
CHANGE	18,218	19,250	1,032	20,950	1,700
NEXT	31,074	32,308	1,234	37,797	5,489
INPUT	88,917	91,619	2,702	102,702	11,083
DELETE	23,423	24,271	848	27,919	3,648
TOP	23,439	24,299	860	28,047	3,748
FILE	63,716	64,372	656	65,175	803
ERASE	14,998	15,695	697	16,953	1,258
EXEC2	62,575	64,136	1,561	66,988	2,852
QUERY	18,990	19,916	926	25,717	5,801
FILEDEF	7,039	7,076	37	8,490	1,414
PRINT	48,912	53,210	4,298	52,337	-873
ASSEMBLE	369,054	363,486	-5,568	372,246	8,760
LOAD	17,350	17,679	329	18,948	1,269
GENMOD	14,815	14,895	80	16,231	1,336
BR14	12,328	12,835	507	14,557	1,722
REXREX	66,266	68,844	2,578	70,614	1,770
TOTAL	1,203,285	1,195,663	-7,622	1,271,695	76,032
AVERAGE	48,131	47,827	-305	50,868	3,041

Table 42. EDF Pathlength

A.1.4.2 SFS User Pathlength Comparisons

The following table contains the pathlength data for the set of commands where the files being used are in the Shared File System. Included here is the virtual pathlength executed in the *user's* virtual machine.

The UXEDIT delta contains significant CRR first time processing. This processing is not unique to UXEDIT. UXEDIT just happens to be the first command in this test sequence where CRR makes a connection to the Recovery Server. Note that the delta for UXEDIT2 is not as large because the first time processing is not present. The UASSEMBLE increase was primarily due to an SP6 APAR fix that resulted in additional SFS server calls.

The average pathlength for the trace increased, but multi-user runs (see "CMS Regression" on page 84) showed a decrease in virtual processor usage (VIRT/CMD under CPU Usage). This is because:

- Several performance improvements were made to CMS functions that are not represented in the traces.
- Nearly all I/O is to SFS for the trace measurements.
- Trace measurements are weighted more heavily with first time processing.

The XA mode data shows the same trend as the EDF data shown above with a 4% increase in virtual pathlength. As described in " EDF Pathlength Comparisons" on page 141, there is a flat increase associated with going from 370 to XA mode. Since the base pathlength of SFS commands is greater than EDF, the percent increase appears smaller. Also, there are fewer fullscreen commands in the SFS sequence.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
UACCESS	37,924	32,857	-5,067	34,731	1,874
UCOPY	19,246	19,462	216	20,883	1,421
URENAME	22,045	22,648	603	24,303	1,655
UXEDIT	75,714	93,474	17,760	98,456	4,982
UFILE	67,469	71,597	4,128	72,923	1,326
UERASE	17,901	18,404	503	19,903	1,499
UASSEMBLE	377,222	393,146	15,924	402,089	8,943
ULOAD	26,434	29,808	3,374	31,250	1,442
UGENMOD	18,122	21,331	3,209	22,838	1,507
UBR14	18,212	21,308	3,096	22,838	1,530
UCREATE	18,280	18,316	36	19,771	1,455
UGRANT	37,746	39,480	1,734	41,417	1,937
UREVOKE	35,751	37,861	2,110	39,760	1,899
UXEDIT2	46,538	52,795	6,257	56,679	3,884
UFILE2	30,636	34,091	3,455	35,376	1,285
UFILEDEF	7,761	6,881	-880	8,517	1,636
TOTAL	857,001	913,459	56,458	951,734	38,275
AVERAGE	53,563	57,091	3,529	59,483	2,392

Table 43. SFS User Pathlength

A.1.4.3 SFS Server Pathlength Comparisons

The following table contains the pathlength data for the set of commands where the files being used are in the Shared File System. Included here is the virtual pathlength executed in the *server* virtual machine.

SASSEMBLE shows an increase for the same reason described above for UASSEMBLE. SXEDIT does not show a first time increase corresponding to UXEDIT because the UXEDIT increase is associated with a connection to the CRR Recovery Server, not the SFS server. SFILEDEF has a pathlength of 0 because the user side no longer issues a refresh directory SFS request unless FST updates are pending. This is the result of a performance enhancement in ESA 1.0.

The SFS server traces showed no significant pathlength impact when the server was running in XA mode. This is because most of the XA impact is in CMS, which accounts for only a small percentage of the instruction execution in the SFS server. The net slight decrease in pathlength is believed to be due to measurement variation.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
SACCESS	119,780	104,657	-15,123	107,483	2,826
SCOPY	65,578	65,077	-501	64,395	-682
SRENAME	38,464	33,976	-4,488	33,180	-796
SXEDIT	50,359	45,976	-4,383	62,213	16,237
SFILE	78,894	69,799	-9,095	71,977	2,178
SERASE	43,607	46,881	3,274	45,520	-1,361
SASSEMBLE	434,298	538,211	103,913	499,947	-38,264
SLOAD	63,131	65,156	2,025	63,373	-1,783
SGENMOD	51,071	51,845	774	50,066	-1.779
SBR14	16,997	16,020	-977	16,339	319
SCREATE	65,687	68,153	2,466	76,068	7,915
SGRANT	55,396	59,979	4,583	51,387	-8,592
SREVOKE	57,337	56,554	-783	56,803	249
SXEDIT2	45,413	45,957	544	44,978	-979
SFILE2	80,499	73,510	-6,989	74,803	1,293
SFILEDEF	2,604	0	-2,604	0	0
TOTAL	1,269,115	1,341,751	72,636	1,318,532	-23,219
AVERAGE	79,320	83,859	4,539	82,408	-1,451

Table 44. SFS Server Pathlength

A.1.5 Non-Shared Storage Comparison

This section contains the virtual non-shared storage data for the commands traced. The first set of trace data contains commands that are using the EDF File System. The second set of trace data contains the user side of commands that use the Shared File System. The third set of traces contains the SFS server side for the same commands. Also shown in the tables are the deltas for running an ESA 1.0 virtual machine in XA mode, as compared to running it in 370 mode on the 370 Feature.

A.1.5.1 EDF Non-Shared Pages

The following table contains the non-shared storage data for the set of commands where the files being used are in the EDF File System. Included here is the virtual non-shared storage referenced in the *user's* virtual machine.

There was no significant difference in the average non-shared pages required for this release.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
SET	24	24	0	26	2
ACCESS	23	25	2	27	2
COPY	27	32	5	35	3
COMPARE	25	27	2	27	0
RENAME	33	31	-2	33	2
LISTFILE	27	26	-1	28	2
XEDIT	44	42	-2	44	2
XXXX	28	27	-1	29	2
LOCATE	20	19	-1	21	2
CHANGE	20	19	-1	21	2
NEXT	19	18	-1	20	2
INPUT	24	23	-1	24	1
DELETE	19	18	-1	20	2
TOP	18	17	-1	19	2
FILE	36	36	0	37	1
ERASE	31	30	-1	32	2
EXEC2	17	18	1	19	1
QUERY	26	25	-1	26	1
FILEDEF	25	24	-1	27	3
PRINT	30	29	-1	30	1
ASSEMBLE	95	104	9	108	4
LOAD	42	43	1	45	2
GENMOD	35	37	2	39	2
BR14	26	28	2	30	2
REXREX	20	22	2	22	0
TOTAL	734	744	10	789	45
AVERAGE	29	29	0	32	2

Table 45. EDF Non-Shared Pages

A.1.5.2 SFS User Non-Shared Storage Comparison

The following table contains the non-shared storage data for the set of commands where the files being used are in the Shared File System. Included here is the virtual non-shared storage referenced in the *user's* virtual machine.

The UXEDIT and UASSEMBLE deltas are accounted for by the observations made in " SFS User Pathlength Comparisons" on page 143. Note that the UXEDIT2 does not have the first time processing overhead.

The average delta of "2" for the 370 MODE column is primarily due to control blocks to support CRR.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
UACCESS	26	28	2	28	0
UCOPY	26	29	3	29	0
URENAME	32	31	-1	31	0
UXEDIT	49	56	7	57	1
UFILE	34	37	3	38	1
UERASE	31	30	-1	30	0
UASSEMBLE	108	118	10	121	3
ULOAD	39	39	0	39	0
UGENMOD	32	34	2	34	0
UBR14	30	31	1	34	3
UCREATE	26	28	2	28	0
UGRANT	42	44	2	44	0
UREVOKE	36	37	1	37	0
UXEDIT2	48	47	-1	48	1
UFILE2	33	33	0	34	1
UFILEDEF	25	22	-3	23	1
TOTAL	617	644	27	655	11
AVERAGE	39	40	2	41	1

Table 46. SFS User Non-Shared Pages

A.1.5.3 SFS Server Non-Shared Storage Comparison

The following table contains the non-shared storage data for the set of commands where the files being used are in the Shared File System. Included here is the virtual non-shared storage referenced in the *server* virtual machine. This includes storage in the *server* virtual machine that contains control blocks and data areas.

The significance of the non-shared pages for any server machine is less than that of user virtual machines. This is because 1) for some of these data areas there is only one copy per server and 2) for other data areas there is one copy per SFS agent structure. Each agent structure supports 20 end users.

The SASSEMBLE delta corresponds to the UASSEMBLE delta described in " SFS User Pathlength Comparisons" on page 143.

The -24 delta for SFILEDEF is accounted for by the SFILEDEF comments in " SFS Server Pathlength Comparisons" on page 144.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
SACCESS	55	53	-2	53	0
SCOPY	61	69	8	56	-13
SRENAME	45	46	1	43	-3
SXEDIT	51	50	-1	61	11
SFILE	61	66	5	59	-7
SERASE	51	56	5	53	-3
SASSEMBLE	61	88	27	73	-15
SLOAD	54	58	4	55	-3
SGENMOD	52	56	4	53	-3
SBR14	35	39	4	37	-2
SCREATE	56	63	7	64	1
SGRANT	53	61	8	53	-8
SREVOKE	49	55	6	50	-5
SXEDIT	44	51	7	47	-4
SFILE2	62	69	7	60	-9
SFILEDEF	24	0	-24	0	0
TOTAL	814	880	66	817	-63
AVERAGE	51	55	4	51	-4

Table 47. SFS Server Non-Shared Storage

A.1.6 Shared Pages

This section contains the virtual shared storage data for the commands traced. The first set of trace data contains commands that are using the EDF File System. The second set of trace data contains the user side of commands that use the Shared File System. The third set of traces contains the SFS server side for the same commands. Also shown in the tables are the deltas for running an ESA 1.0 virtual machine in XA mode, as compared to running it in 370 mode on the 370 Feature.

The data labeled UNIQUE represents the number of unique shared pages referenced for all of the commands in the table. This is the total number of different shared pages required to execute all of the commands. Unique pages is a better reflection of the change in real storage required to back these virtual pages.

A.1.6.1 EDF Shared Pages

The following table contains the shared storage data for the set of commands where the files being used are in the EDF File System. Included here is the virtual shared storage referenced in the *user's* virtual machine.

The increase of seven unique shared pages reflects growth in existing shared modules as well as new modules to support the new functions of this release. This increase is not significant enough to cause any noticeable performance differences in this release. The XA mode shared storage requirements are within a few pages of the 370 mode data and again have little overall effect.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
SET	41	40	-1	42	2
ACCESS	53	57	4	61	4
COPY	44	49	5	51	2
COMPARE	35	43	8	44	1
RENAME	46	49	3	52	2
LISTFILE	37	39	2	42	3
XEDIT	72	81	9	82	1
XXXX	63	66	3	73	6
LOCATE	41	42	1	45	3
CHANGE	52	53	1	58	5
NEXT	40	40	0	44	4
INPUT	57	58	1	61	3
DELETE	42	43	1	47	4
TOP	40	41	1	45	4
FILE	67	73	6	76	3
ERASE	43	47	4	49	2
EXEC2	33	37	4	38	1
QUERY	47	46	-1	51	5
FILEDEF	30	32	2	34	2
PRINT	49	54	5	55	1
ASSEMBLE	65	71	6	72	1
LOAD	48	53	5	54	1
GENMOD	44	47	3	48	1
BR14	39	48	9	50	2
REXREX	38	45	7	46	1
TOTAL	1,166	1,254	88	1,320	66
AVERAGE	47	50	4	53	3
UNIQUE	206	213	7	219	6

Table 48. EDF Shared Pages

A.1.6.2 SFS User Shared Storage Comparison

The following table contains the shared storage data for the set of commands where the files being used are in the Shared File System. Included here is the virtual shared storage referenced in the *user's* virtual machine.

The UXEDIT and UASSEMBLE deltas are accounted for by the observations made in " SFS User Pathlength Comparisons" on page 143. UFILE and UFILE2 are larger mostly because of the CRR registration/unregistration for a non-system work unit.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
UACCESS	65	74	9	75	1
UCOPY	56	64	8	64	0
URENAME	60	67	7	67	0
UXEDIT	98	125	27	125	0
UFILE	86	112	26	114	2
UERASE	62	73	11	72	-1
UASSEMBLE	96	127	31	127	0
ULOAD	64	86	22	85	-1
UGENMOD	59	81	22	80	-1
UBR14	55	80	25	80	0
UCREATE	61	74	13	73	-1
UGRANT	84	95	11	94	-1
UREVOKE	80	94	14	93	-1
UXEDIT2	97	114	17	114	0
UFILE2	85	112	27	114	2
UFILEDEF	36	32	-4	33	1
TOTAL	1,144	1,410	266	1,410	0
AVERAGE	72	88	16	88	0
UNIQUE	207	241	34	242	1

Table 49. SFS User Shared Pages

A.1.6.3 SFS Server Shared Storage Comparison

The following table contains the shared storage data for the set of commands where the files being used are in the Shared File System. Included here is the shared virtual storage referenced in the *server* virtual machine. In this context, shared storage refers to code running from a shared saved segment. For this section, the data was divided into two sets of data because of the different usage of the segments. This code includes both SFS server code (DMSDAC and DMSSAC) and CMS. The CMS code referenced in the SFS server is shared with other CMS users. The DAC/SAC code would be shared only with other SFS servers on the same system.

The SASSEMBLE delta corresponds to the delta described for UASSEMBLE in “ SFS User Pathlength Comparisons” on page 143.

	370 MODE						XA MODE			
COMMAND	SP6		ESA 1.0		DELTA		ESA 1.0		DELTA	
	DAC/ SAC	CMS	DAC/ SAC	CMS	DAC/ SAC	CMS	DAC/ SAC	CMS	DAC/ SAC	CMS
SACCESS	44	2	47	3	3	1	51	3	4	0
SCOPY	53	6	60	11	7	5	61	9	1	-2
SRENAME	49	2	48	3	-1	1	48	7	0	4
SXEDIT	55	5	60	3	5	-2	70	19	10	16
SFILE	63	5	69	9	6	4	69	11	0	2
SERASE	53	2	59	3	6	1	59	7	0	4
SASSEMBLE	62	2	81	15	19	13	82	17	1	2
SLOAD	50	2	58	3	8	1	58	7	0	4
SGENMOD	50	2	54	3	4	1	54	7	0	4
SBR14	39	2	44	3	5	1	44	7	0	4
SCREATE	49	2	53	3	4	1	67	13	14	10
SGRANT	54	2	56	3	2	1	57	7	1	4
SREVOKE	54	2	59	3	5	1	59	7	0	4
SXEDIT2	55	2	60	3	5	1	60	7	0	4
SFILE2	63	5	69	9	6	4	69	11	0	2
SFILEDEF	14	2	0	0	-14	-2	0	0	0	0
TOTAL	807	45	877	77	70	32	908	139	31	62
AVERAGE	50	3	55	5	5	2	57	9	2	4
UNIQUE	92	6	119	15	27	9	122	19	3	4

Table 50. SFS Server Shared Storage

A.1.7 Special Operation Comparison

This section contains the virtual special operations data for the commands traced. Special operations are those instructions which cause either CP or the hardware microcode to provide services for the virtual machine. This includes privileged instructions, normally invalid opcodes (i.e. DIAG), and the first time a non-shared page is referenced. The first set of trace data contains commands that are using the EDF File System. The second set of trace data contains the user side of commands that use the Shared File System. The third set of traces contains the SFS server side for the same commands. Also shown in the tables are the deltas for running an ESA 1.0 virtual machine in XA mode, as compared to running it in 370 mode on the 370 Feature.

A.1.7.1 EDF Special Operation Comparison

The following table contains the special operations data for the set of commands where the files being used are in the EDF File System. Included here is the virtual special operations executed in the *user's* virtual machine.

There was no significant difference in the average special operations for the 370 mode comparison. The XA mode special operation data shows an increase over the 370 mode data. This is mostly due to the addition of STNSM/STOSM and LCTL/STCTL instructions to support this new environment.

Also traced, but not shown in the table, was a REXX EXEC which makes heavy usage of EXECIO. Improvements were made in this area, and tracing was used to verify and quantify the benefit. The special operations were reduced from 372 in SP6 to 164 in ESA 1.0 (a 56% reduction). This improvement is due to reduced overhead in the linkage between REXX and EXECIO by eliminating SPKA, LPSW, and SVC instructions.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
SET	45	45	0	89	44
ACCESS	79	81	2	170	89
COPY	80	86	6	139	53
COMPARE	255	257	2	597	340
RENAME	62	60	-2	109	49
LISTFILE	58	55	-3	128	73
XEDIT	99	105	6	208	103
XXXX	72	66	-6	128	62
LOCATE	52	46	-6	92	46
CHANGE	46	47	1	93	46
NEXT	51	45	-6	91	46
INPUT	100	86	-14	179	97
DELETE	51	45	-6	91	46
TOP	50	44	-6	90	46
FILE	106	113	7	184	71
ERASE	60	59	-1	108	49
EXEC2	292	293	1	721	428
QUERY	103	102	-1	286	184
FILEDEF	46	45	-1	95	50
PRINT	72	71	-1	113	42
ASSEMBLE	1,012	1,025	13	1,875	850
LOAD	88	90	2	140	60
GENMOD	78	80	2	134	54
BR14	58	62	4	111	49
REXREX	106	110	4	252	142
TOTAL	3,121	3,118	-3	6,223	3,105
AVERAGE	125	125	0	249	124

Table 51. EDF Special Operations

A.1.7.2 SFS User Special Operations

The following table contains the special operations data for the set of commands where the files being used are in the Shared File System. Included here is the virtual special operations executed in the *user's* virtual machine.

The UXEDIT and UASSEMBLE deltas are accounted for by the observations made in " SFS User Pathlength Comparisons" on page 143.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
UACCESS	75	74	-1	143	69
UCOPY	54	57	3	102	45
URENAME	57	55	-2	102	47
UXEDIT	114	146	32	250	104
UFILE	266	291	25	362	71
UERASE	56	54	-2	101	47
UASSEMBLE	1,037	1,065	28	1,881	816
ULOAD	89	97	8	145	48
UGENMOD	71	82	11	134	52
UBR14	69	81	12	134	53
UCREATE	51	52	1	99	47
UGRANT	101	102	1	212	110
UREVOKE	88	88	0	180	92
UXEDIT2	114	129	15	237	108
UFILE2	89	112	23	182	70
UFILEDEF	47	43	-4	91	48
TOTAL	2,378	2,528	150	4,355	1,827
AVERAGE	149	158	9	272	114

Table 52. SFS User Special Operations

A.1.7.3 SFS Server Special Operations

The following table contains the special operations data for the set of commands where the files being used are in the Shared File System. Included here is the virtual special operations executed in the *server* virtual machine.

The improvements from SP 6 are due to one fewer refresh directory server call in each command and to enables/disables being removed from the SFS dispatcher scan.

	370 MODE			XA MODE	
COMMAND	SP6	ESA 1.0	DELTA	ESA 1.0	DELTA
SACCESS	188	71	-117	271	200
SCOPY	129	132	3	263	131
SRENAME	85	72	-13	124	52
SXEDIT	135	106	-29	375	269
SFILE	170	143	-27	344	201
SERASE	91	82	-9	134	52
SASSEMBLE	867	822	-45	2,189	1,367
SLOAD	158	137	-21	296	159
SGENMOD	117	103	-14	196	93
SBR14	81	69	-12	131	62
SCREATE	96	89	-7	214	125
SGRANT	102	95	-7	134	39
SREVOKE	89	81	-8	131	50
SXEDIT2	119	107	-12	215	108
SFILE2	170	162	-8	345	183
SFILEDEF	31	0	-31	0	0
TOTAL	2,628	2,271	-357	5,362	3,091
AVERAGE	164	142	-22	335	193

Table 53. SFS Server Special Operations

Appendix B. SFS Counter Data

B.1 Introduction

The following SFS counts and timings are provided to supplement the SFS measurements. These were acquired by issuing the QUERY FILEPOOL STATUS command once at the beginning of the measurement interval and once at the end.

For the large system measurements, QUERY FILEPOOL STATUS (QFS) data was obtained for each SFS file pool and the counts and timings for each file pool were added together.

The QFS data was also collected for the CRR Recovery Server, but that data is not included.

The first section in each table consists of the counters normalized by the number of commands. The beginning values were subtracted from the ending values and divided by the number of commands in the interval. For each table, counts and timings which have all zero values are typically not shown. A description of the SFS counts and timings can be found in the ***VM/ESA 1.0 CMS Administration Reference*** book.

The second section in each table consists of derived relationships which were calculated from a combination of two or more individual counts and/or timings. See “Glossary of Performance Terms” on page 194 for definitions of these derived values.

B.1.1 3090-600J

SECTION	3.4.2.1
RUN ID	Y63F4804
PROCESSOR	3090-600J
REAL STORAGE	512M
EXP. STORAGE	2048M
WORKLOAD	FS7B35
USERS	4800
NORMALIZED BY COMMAND	
Checkpoints Taken	0.001
Checkpoint Time (msec)	3.580
Close File Requests	0.397
Connect Requests	0.006
Delete File Requests	0.095
Lock Requests	0.025
Open File New Requests	0.002
Open File Read Requests	0.227
Open File Replace Requests	0.145
Open File Write Requests	0.023
Query File Pool Requests	0.000
Query User Space Requests	0.021
Read File Requests	0.229
Refresh Directory Requests	0.013
Rename Requests	0.005
Unlock Requests	0.025
Write File Requests	0.128
Total File Pool Requests	1.341
File Pool Request Service Time (msec)	150.268
Local File Pool Requests	1.341
Begin LUWs	0.507
Agent Holding Time (msec)	188.025
LUW Rollbacks	0.000
SAC Calls	6.244
Catalog Lock Conflicts	0.008
Total Lock Conflicts	0.008
Lock Wait Time (msec)	0.668
Deadlocks	0.000
File Blocks Read	0.949
File Blocks Written	0.587
Catalog Blocks Read	0.436
Catalog Blocks Written	0.282
Ctl Minidisk Blocks Read	0.000
Ctl Minidisk Blocks Written	0.113
Log Blocks Written	0.442
Total DASD Block Transfers	2.809
BIO Requests to Read File Blocks	0.506
BIO Requests to Write File Blocks	0.293
BIO Requests to Read Catalog Blocks	0.436
BIO Requests to Write Catalog Blocks	0.282
BIO Requests to Read Ctl Minidisk Blocks	0.000
BIO Requests to Write Ctl Minidisk Blocks	0.004
BIO Requests to Write Log Blocks	0.442
Total BIO Requests	1.963
Total BIO Request Time (msec)	45.425
I/O Requests to Read File Blocks	0.516
I/O Requests to Write File Blocks	0.381
I/O Requests to Read Catalog Blocks	0.436
I/O Requests to Write Catalog Blocks	0.282
I/O Requests to Read Ctl Minidisk Blocks	0.000
I/O Requests to Write Ctl Minidisk Blocks	0.007
I/O Requests to Write Log Blocks	0.442
Total I/O Requests	2.064
DERIVED RELATIONSHIPS - PRODUCTION FILEPOOLS ONLY	
Deadlocks	4
Rollbacks Due to Deadlock	0
LUW Rollbacks	1
Checkpoints Taken	203
Checkpoint Duration (sec)	5.56
Seconds Between Checkpoints	37.19
Checkpoint Utilization	14.94
Agents Held	7.84
Agents In-call	6.27
MSec per BIO Request	23.14
Blocking Factor (Blocks/BIO)	1.43
Chaining Factor (Blocks/BIO)	1.36

B.1.2 4381-13 / 35% SFS

SECTIONS	4.2.2.2 4.7.2.1	4.2.2.2 4.7.2.1	4.2.2.2 4.7.2.1	4.2.2.2 4.3.2.1 4.7.2.1	4.2.2.2 4.3.2.1 4.7.2.1	4.2.2.2 4.3.2.1 4.7.2.1
RELEASE	SP6	SP6	SP6	ESA 1.0	ESA 1.0	ESA 1.0
RUN ID	EC4292	EC4292	EC4292	EC7600	EC7600	EC7600
PROCESSOR	4381-13	4381-13	4381-13	4381-13	4381-13	4381-13
STORAGE	16M	16M	16M	16M	16M	16M
WORKLOAD	35% SFS	35% SFS	35% SFS	35% SFS	35% SFS	35% SFS
ACTIVE USERS	59	67	75	59	67	75
NORMALIZED BY COMMAND						
Checkpoints Taken	0.001	0.001	0.001	0.001	0.001	0.001
Checkpoint Time (msec)	1.626	1.713	1.807	1.664	1.995	2.184
Close File Requests	0.438	0.435	0.434	0.452	0.439	0.434
Commit Requests	0.013	0.012	0.012	0.000	0.000	0.000
Connect Requests	0.000	0.000	0.000	0.003	0.003	0.003
Delete File Requests	0.111	0.110	0.111	0.108	0.108	0.105
Get Directory Entry Requests	0.004	0.003	0.004	0.000	0.000	0.000
Lock Requests	0.027	0.027	0.027	0.027	0.026	0.027
Open File New Requests	0.002	0.002	0.002	0.002	0.002	0.002
Open File Read Requests	0.251	0.248	0.248	0.260	0.247	0.247
Open File Replace Requests	0.152	0.152	0.153	0.165	0.165	0.159
Open File Write Requests	0.032	0.033	0.032	0.025	0.025	0.025
Query File Pool Requests	0.000	0.000	0.000	0.000	0.000	0.000
Query User Space Requests	0.023	0.023	0.023	0.023	0.024	0.023
Read File Requests	0.290	0.290	0.285	0.264	0.260	0.254
Refresh Directory Requests	0.575	0.575	0.577	0.014	0.014	0.014
Rename Requests	0.005	0.005	0.005	0.005	0.005	0.005
Unlock Requests	0.026	0.026	0.027	0.027	0.027	0.026
Write File Requests	0.158	0.155	0.156	0.148	0.148	0.141
Total File Pool Requests	2.108	2.095	2.096	1.524	1.492	1.466
File Pool Request Service Time (msec)	104.052	125.389	160.768	95.643	117.522	144.439
Local File Pool Requests	2.108	2.095	2.096	1.524	1.492	1.466
Begin LUWs	0.538	0.533	0.537	0.565	0.546	0.542
Agent Holding Time (msec)	334.981	410.086	519.974	288.833	363.632	452.053
SAC Calls	6.531	6.480	6.503	7.101	6.886	6.733
Catalog Lock Conflicts	0.001	0.001	0.004	0.001	0.002	0.004
Total Lock Conflicts	0.001	0.001	0.004	0.001	0.002	0.004
Lock Wait Time (msec)	0.018	0.013	0.220	0.021	0.063	0.110
File Blocks Read	1.116	1.112	1.097	1.088	1.063	1.046
File Blocks Written	0.695	0.684	0.685	0.669	0.670	0.643
Catalog Blocks Read	0.580	0.646	0.690	0.630	0.618	0.657
Catalog Blocks Written	0.255	0.276	0.300	0.288	0.291	0.307
Ctl Minidisk Blocks Written	0.023	0.021	0.020	0.022	0.024	0.019
Log Blocks Written	0.587	0.578	0.577	0.623	0.606	0.592
Total DASD Block Transfers	3.257	3.317	3.369	3.320	3.271	3.263
BIO Requests to Read File Blocks	0.577	0.580	0.584	0.559	0.541	0.540
BIO Requests to Write File Blocks	0.339	0.325	0.330	0.320	0.322	0.310
BIO Requests to Read Catalog Blocks	0.580	0.646	0.690	0.630	0.618	0.657
BIO Requests to Write Catalog Blocks	0.255	0.276	0.300	0.288	0.291	0.307
BIO Requests to Write Ctl Minidisk Block	0.005	0.004	0.004	0.005	0.005	0.004
BIO Requests to Write Log Blocks	0.587	0.578	0.577	0.623	0.606	0.592
Total BIO Requests	2.334	2.410	2.485	2.425	2.383	2.410
Total BIO Request Time (msec)	67.070	83.085	110.464	66.399	79.932	97.751
I/O Requests to Read File Blocks	0.587	0.592	0.593	0.561	0.543	0.541
I/O Requests to Write File Blocks	0.339	0.334	0.338	0.320	0.322	0.310
I/O Requests to Read Catalog Blocks	0.580	0.646	0.690	0.630	0.618	0.657
I/O Requests to Write Catalog Blocks	0.255	0.276	0.300	0.288	0.291	0.307
I/O Requests to Write Ctl Minidisk Block	0.007	0.006	0.006	0.007	0.007	0.006
I/O Requests to Write Log Blocks	0.587	0.578	0.577	0.623	0.606	0.592
Total I/O Requests	2.355	2.432	2.503	2.428	2.387	2.413
DERIVED RELATIONSHIPS						
Deadlocks	0	0	0	0	0	0
Rollbacks Due to Deadlock	0	0	0	0	0	0
LUW Rollbacks	0	0	0	0	0	0
Checkpoints Taken	5	5	5	5	6	5
Checkpoint Duration (sec)	2.11	2.48	2.82	2.19	2.46	3.50
Seconds Between Checkpoints	240.60	240.80	240.69	241.00	200.83	242.40
Checkpoint Utilization	0.88	1.03	1.17	0.91	1.22	1.45
Agents Held	1.81	2.47	3.37	1.58	2.23	2.99
Agents In-call	0.56	0.75	1.04	0.52	0.72	0.96
Msec per BIO Request	28.74	34.48	44.46	27.38	33.54	40.57
Blocking Factor (Blocks/BIO)	1.40	1.38	1.36	1.37	1.37	1.35
Chaining Factor (Blocks/IO)	1.38	1.36	1.35	1.37	1.37	1.35

B.1.3 4381-13 / MAX% SFS

SECTION	4.2.2.3	4.2.2.3	4.2.2.3	4.2.2.3	4.2.2.3	4.2.2.3
RELEASE	SP 6	SP 6	SP 6	ESA 1.0	ESA 1.0	ESA 1.0
RUN ID	EC4599	EC4599	EC4599	EC7604	EC7604	EC7604
PROCESSOR	4381-13	4381-13	4381-13	4381-13	4381-13	4381-13
STORAGE	16M	16M	16M	16M	16M	16M
WORKLOAD	MAX% SFS	MAX% SFS	MAX% SFS	MAX% SFS	MAX% SFS	MAX% SFS
ACTIVE USERS	57	65	73	57	65	73
NORMALIZED BY COMMAND						
Checkpoints Taken	0.001	0.001	0.001	0.001	0.001	0.001
Checkpoint Time (msec)	1.037	1.495	1.750	1.051	1.280	1.344
Close File Requests	0.747	0.772	0.765	0.766	0.741	0.799
Commit Requests	0.026	0.026	0.025	0.000	0.000	0.000
Connect Requests	0.000	0.000	0.000	0.005	0.005	0.006
Delete File Requests	0.112	0.110	0.110	0.103	0.106	0.105
File Copy Requests	0.003	0.003	0.004	0.004	0.003	0.004
Get Directory Entry Requests	0.003	0.004	0.004	0.000	0.000	0.000
Lock Requests	0.026	0.027	0.027	0.027	0.027	0.027
Open File Read Requests	0.563	0.591	0.585	0.584	0.557	0.608
Open File Replace Requests	0.151	0.149	0.149	0.158	0.160	0.157
Open File Write Requests	0.032	0.032	0.032	0.025	0.025	0.025
Query File Pool Requests	0.000	0.000	0.000	0.000	0.000	0.000
Query User Space Requests	0.024	0.023	0.023	0.023	0.023	0.023
Read File Requests	0.416	0.411	0.416	0.384	0.378	0.388
Refresh Directory Requests	0.602	0.602	0.604	0.028	0.028	0.029
Rename Requests	0.005	0.005	0.005	0.005	0.006	0.005
Unlock Requests	0.026	0.027	0.027	0.027	0.027	0.026
Write File Requests	0.160	0.154	0.157	0.142	0.142	0.140
Total File Pool Requests	2.897	2.936	2.935	2.282	2.227	2.333
File Pool Request Service Time (msec)	126.139	162.601	197.391	129.224	145.226	187.510
Local File Pool Requests	2.897	2.936	2.935	2.282	2.227	2.333
Begin LUWs	0.833	0.863	0.857	0.864	0.836	0.888
Agent Holding Time (msec)	369.351	495.557	631.408	362.519	422.266	583.282
LUW Rollbacks	0.000	0.000	0.000	0.000	0.000	0.000
SAC Calls	9.526	9.699	9.666	10.044	9.789	10.166
Catalog Lock Conflicts	0.002	0.001	0.004	0.002	0.002	0.004
Total Lock Conflicts	0.002	0.001	0.004	0.002	0.002	0.004
Lock Wait Time (msec)	0.022	0.003	0.322	0.045	0.024	0.189
Deadlocks	0.000	0.000	0.000	0.000	0.000	0.000
File Blocks Read	1.759	1.772	1.784	1.725	1.680	1.769
File Blocks Written	0.698	0.679	0.686	0.643	0.647	0.637
Catalog Blocks Read	1.160	1.207	1.252	1.203	1.272	1.368
Catalog Blocks Written	0.280	0.317	0.333	0.322	0.349	0.351
Ctl Minidisk Blocks Written	0.019	0.021	0.020	0.018	0.020	0.019
Log Blocks Written	0.585	0.579	0.581	0.607	0.605	0.597
Total DASD Block Transfers	4.501	4.575	4.657	4.518	4.574	4.722
BIO Requests to Read File Blocks	1.022	1.057	1.051	1.014	0.977	1.043
BIO Requests to Write File Blocks	0.332	0.325	0.328	0.308	0.310	0.308
BIO Requests to Read Catalog Blocks	1.160	1.207	1.252	1.203	1.272	1.368
BIO Requests to Write Catalog Blocks	0.280	0.317	0.333	0.322	0.349	0.351
BIO Requests to Write Ctl Minidisk Block	0.004	0.004	0.004	0.004	0.004	0.004
BIO Requests to Write Log Blocks	0.585	0.579	0.581	0.607	0.605	0.597
Total BIO Requests	3.382	3.489	3.548	3.458	3.517	3.661
Total BIO Request Time (msec)	83.528	110.556	133.064	88.967	102.404	136.630
I/O Requests to Read File Blocks	1.027	1.065	1.057	1.015	0.978	1.043
I/O Requests to Write File Blocks	0.338	0.333	0.334	0.309	0.310	0.308
I/O Requests to Read Catalog Blocks	1.160	1.207	1.252	1.203	1.272	1.368
I/O Requests to Write Catalog Blocks	0.280	0.317	0.333	0.322	0.349	0.351
I/O Requests to Write Ctl Minidisk Block	0.006	0.006	0.006	0.006	0.006	0.006
I/O Requests to Write Log Blocks	0.585	0.579	0.581	0.607	0.605	0.597
Total I/O Requests	3.396	3.507	3.562	3.461	3.520	3.664
DERIVED RELATIONSHIPS						
Deadlocks	0	0	1	0	0	0
Rollbacks Due to Deadlock	0	0	0	0	0	0
LUW Rollbacks	0	0	1	0	0	0
Checkpoints Taken	4	5	5	4	5	5
Checkpoint Duration (sec)	1.63	2.09	2.72	1.67	1.83	2.08
Seconds Between Checkpoints	300.50	240.40	240.80	301.75	241.20	242.20
Checkpoint Utilization	0.54	0.87	1.13	0.55	0.76	0.86
Agents Held	1.94	2.89	4.07	1.91	2.51	3.73
Agents In-call	0.66	0.95	1.27	0.68	0.86	1.20
Msec per BIO Request	24.69	31.69	37.50	25.73	29.12	37.32
Blocking Factor (Blocks/BIO)	1.33	1.31	1.31	1.31	1.30	1.29
Chaining Factor (Blocks/IO)	1.33	1.30	1.31	1.31	1.30	1.29

B.1.4 9370-80

SECTION	4.2.2.4	4.2.2.4	4.2.2.4	4.2.2.4	4.2.2.4	4.2.2.4
RELEASE	SP6	SP6	SP6	ESA1.0	ESA1.0	ESA1.0
RUN ID	EC7601	EC7601	EC7601	EC7033	EC7033	EC7033
PROCESSOR	9370-80	9370-80	9370-80	9370-80	9370-80	9370-80
STORAGE	16M	16M	16M	16M	16M	16M
WORKLOAD	35% SFS	35% SFS	35% SFS	35% SFS	35% SFS	35% SFS
ACTIVE USERS	27	31	35	27	31	35
NORMALIZED BY COMMAND						
Checkpoints Taken	0.001	0.001	0.001	0.001	0.001	0.001
Checkpoint Time (msec)	0.900	1.037	1.284	1.023	1.537	1.910
Close File Requests	0.436	0.436	0.435	0.440	0.435	0.436
Commit Requests	0.012	0.012	0.012	0.000	0.000	0.000
Connect Requests	0.000	0.000	0.000	0.003	0.004	0.004
Delete File Requests	0.112	0.111	0.111	0.106	0.105	0.106
Get Directory Entry Requests	0.004	0.004	0.004	0.000	0.000	0.000
Lock Requests	0.027	0.026	0.027	0.027	0.027	0.027
Open File New Requests	0.002	0.002	0.002	0.002	0.002	0.002
Open File Read Requests	0.249	0.249	0.248	0.251	0.247	0.248
Open File Replace Requests	0.153	0.153	0.152	0.162	0.162	0.161
Open File Write Requests	0.033	0.032	0.032	0.025	0.025	0.025
Query File Pool Requests	0.000	0.000	0.000	0.000	0.000	0.000
Query User Space Requests	0.023	0.023	0.023	0.023	0.023	0.023
Read File Requests	0.285	0.289	0.283	0.260	0.254	0.254
Refresh Directory Requests	0.573	0.576	0.576	0.013	0.014	0.014
Rename Requests	0.005	0.005	0.005	0.005	0.005	0.005
Unlock Requests	0.027	0.027	0.027	0.027	0.027	0.027
Write File Requests	0.155	0.157	0.153	0.145	0.142	0.143
Total File Pool Requests	2.095	2.103	2.090	1.491	1.471	1.476
File Pool Request Service Time (msec)	130.840	144.622	173.626	119.491	132.265	148.312
Local File Pool Requests	2.095	2.103	2.090	1.491	1.471	1.476
Begin LUWs	0.537	0.537	0.536	0.559	0.545	0.545
Agent Holding Time (msec)	523.005	619.909	874.458	517.775	591.405	772.292
SAC Calls	6.524	6.519	6.492	6.884	6.799	6.748
Catalog Lock Conflicts	0.000	0.002	0.002	0.001	0.001	0.000
Total Lock Conflicts	0.000	0.002	0.002	0.001	0.001	0.000
Lock Wait Time (msec)	0.002	0.074	0.110	0.013	0.022	0.010
File Blocks Read	1.099	1.108	1.093	1.073	1.044	1.048
File Blocks Written	0.686	0.691	0.679	0.656	0.648	0.649
Catalog Blocks Read	0.630	0.694	0.781	0.656	0.706	0.764
Catalog Blocks Written	0.255	0.277	0.314	0.283	0.291	0.319
Ctl Minidisk Blocks Written	0.019	0.017	0.016	0.018	0.020	0.018
Log Blocks Written	0.591	0.586	0.585	0.612	0.604	0.600
Total DASD Block Transfers	3.280	3.373	3.468	3.297	3.313	3.398
BIO Requests to Read File Blocks	0.576	0.579	0.578	0.549	0.540	0.544
BIO Requests to Write File Blocks	0.327	0.328	0.325	0.316	0.309	0.312
BIO Requests to Read Catalog Blocks	0.630	0.694	0.781	0.656	0.706	0.764
BIO Requests to Write Catalog Blocks	0.255	0.277	0.314	0.283	0.291	0.319
BIO Requests to Write Ctl Minidisk Block	0.004	0.004	0.003	0.004	0.004	0.004
BIO Requests to Write Log Blocks	0.591	0.586	0.585	0.612	0.604	0.600
Total BIO Requests	2.384	2.469	2.587	2.420	2.455	2.543
Total BIO Request Time (msec)	82.798	93.276	113.766	86.895	94.656	106.619
I/O Requests to Read File Blocks	0.608	0.620	0.610	0.569	0.556	0.559
I/O Requests to Write File Blocks	0.368	0.371	0.358	0.328	0.320	0.323
I/O Requests to Read Catalog Blocks	0.630	0.694	0.781	0.656	0.706	0.764
I/O Requests to Write Catalog Blocks	0.255	0.277	0.314	0.283	0.291	0.319
I/O Requests to Write Ctl Minidisk Block	0.007	0.006	0.006	0.007	0.008	0.007
I/O Requests to Write Log Blocks	0.591	0.586	0.585	0.612	0.604	0.600
Total I/O Requests	2.460	2.554	2.654	2.456	2.486	2.572
DERIVED RELATIONSHIPS						
Deadlocks	0	0	0	0	0	0
Rollbacks Due to Deadlock	0	0	0	0	0	0
LUW Rollbacks	0	0	0	0	0	0
Checkpoints Taken	5	5	5	5	6	6
Checkpoint Duration (sec)	1.32	1.72	2.30	1.52	2.12	2.87
Seconds Between Checkpoints	600.60	600.60	600.60	601.40	501.33	501.50
Checkpoint Utilization	0.22	0.29	0.38	0.25	0.42	0.57
Agents Held	1.27	1.71	2.61	1.28	1.63	2.31
Agents In-call	0.32	0.40	0.52	0.30	0.36	0.44
MSec per BIO Request	34.73	37.79	43.98	35.90	38.56	41.93
Blocking Factor (Blocks/BIO)	1.38	1.37	1.34	1.36	1.36	1.34
Chaining Factor (Blocks/IO)	1.33	1.32	1.31	1.34	1.33	1.32

B.1.5 9370-30

SECTION	4.2.2.5	4.2.2.5	4.2.2.5	4.2.2.5	4.2.2.5	4.2.2.5
RELEASE	SP6	SP6	SP6	ESA1.0	ESA1.0	ESA1.0
RUN ID	E05437	E05438	E05440	EC7606	EC7606	EC7606
PROCESSOR	9370-30	9370-30	9370-30	9370-30	9370-30	9370-30
STORAGE	8M	8M	8M	8M	8M	8M
WORKLOAD	35% SFS	35% SFS	35% SFS	35% SFS	35% SFS	35% SFS
ACTIVE USERS	17	19	21	17	19	21
NORMALIZED BY COMMAND						
Checkpoints Taken	0.001	0.001	0.001	0.001	0.001	0.001
Checkpoint Time (msec)	1.770	1.600	1.861	2.125	2.311	1.629
Close File Requests	0.441	0.431	0.430	0.438	0.442	0.438
Commit Requests	0.012	0.013	0.013	0.000	0.000	0.000
Connect Requests	0.000	0.000	0.000	0.004	0.004	0.004
Delete File Requests	0.113	0.111	0.111	0.106	0.108	0.105
Get Directory Entry Requests	0.003	0.003	0.004	0.000	0.000	0.000
Lock Requests	0.027	0.027	0.027	0.027	0.027	0.027
Open File New Requests	0.002	0.002	0.002	0.002	0.002	0.002
Open File Read Requests	0.254	0.245	0.245	0.251	0.251	0.251
Open File Replace Requests	0.153	0.152	0.151	0.160	0.162	0.161
Open File Write Requests	0.032	0.033	0.032	0.025	0.025	0.025
Query File Pool Requests	0.000	0.000	0.000	0.000	0.000	0.000
Query User Space Requests	0.023	0.023	0.023	0.023	0.023	0.023
Read File Requests	0.287	0.291	0.288	0.251	0.257	0.253
Refresh Directory Requests	0.577	0.577	0.575	0.014	0.014	0.014
Rename Requests	0.005	0.005	0.005	0.006	0.005	0.005
Unlock Requests	0.026	0.027	0.027	0.027	0.027	0.027
Write File Requests	0.154	0.158	0.156	0.140	0.146	0.139
Total File Pool Requests	2.110	2.098	2.090	1.472	1.492	1.474
File Pool Request Service Time (msec)	178.731	194.319	211.874	165.872	180.480	199.306
Local File Pool Requests	2.110	2.098	2.090	1.472	1.492	1.474
Begin LUWs	0.543	0.533	0.534	0.551	0.548	0.545
Agent Holding Time (msec)	940.706	1035.746	1210.837	779.670	913.873	1041.281
SAC Calls	6.571	6.484	6.478	6.833	6.793	6.727
Catalog Lock Conflicts	0.001	0.001	0.002	0.002	0.001	0.001
Total Lock Conflicts	0.001	0.001	0.002	0.002	0.001	0.001
Lock Wait Time (msec)	0.025	0.128	0.095	0.070	0.043	0.079
File Blocks Read	1.107	1.115	1.107	1.038	1.058	1.046
File Blocks Written	0.682	0.691	0.683	0.640	0.659	0.636
Catalog Blocks Read	0.691	0.728	0.808	0.712	0.735	0.769
Catalog Blocks Written	0.270	0.278	0.315	0.294	0.305	0.315
Ctl Minidisk Blocks Read	0.000	0.000	0.000	0.000	0.000	0.000
Ctl Minidisk Blocks Written	0.019	0.018	0.017	0.019	0.021	0.016
Log Blocks Written	0.593	0.586	0.586	0.608	0.604	0.593
Total DASD Block Transfers	3.363	3.416	3.515	3.310	3.382	3.376
BIO Requests to Read File Blocks	0.588	0.587	0.582	0.540	0.548	0.548
BIO Requests to Write File Blocks	0.329	0.333	0.328	0.308	0.318	0.307
BIO Requests to Read Catalog Blocks	0.691	0.728	0.808	0.712	0.735	0.769
BIO Requests to Write Catalog Blocks	0.270	0.278	0.315	0.294	0.305	0.315
BIO Requests to Read Ctl Minidisk Blocks	0.000	0.000	0.000	0.000	0.000	0.000
BIO Requests to Write Ctl Minidisk Block	0.004	0.004	0.004	0.004	0.005	0.004
BIO Requests to Write Log Blocks	0.593	0.586	0.586	0.608	0.604	0.593
Total BIO Requests	2.475	2.516	2.623	2.466	2.514	2.536
Total BIO Request Time (msec)	116.012	130.614	145.836	116.021	127.541	141.497
I/O Requests to Read File Blocks	0.621	0.629	0.623	0.561	0.570	0.563
I/O Requests to Write File Blocks	0.371	0.373	0.370	0.323	0.330	0.321
I/O Requests to Read Catalog Blocks	0.691	0.728	0.808	0.712	0.735	0.769
I/O Requests to Write Catalog Blocks	0.270	0.278	0.315	0.294	0.305	0.315
I/O Requests to Read Ctl Minidisk Blocks	0.000	0.000	0.000	0.000	0.000	0.000
I/O Requests to Write Ctl Minidisk Block	0.008	0.007	0.007	0.007	0.008	0.006
I/O Requests to Write Log Blocks	0.593	0.586	0.586	0.608	0.604	0.593
Total I/O Requests	2.565	2.602	2.709	2.505	2.552	2.567
DERIVED RELATIONSHIPS						
Deadlocks	0	0	0	0	0	0
Rollbacks Due to Deadlock	0	0	0	0	0	0
LUW Rollbacks	0	0	0	0	0	0
Checkpoints Taken	5	5	5	5	6	5
Checkpoint Duration (sec)	2.45	2.42	3.03	3.00	2.95	2.68
Seconds Between Checkpoints	960.60	960.60	960.60	961.60	861.83	962.00
Checkpoint Utilization	0.25	0.25	0.32	0.31	0.37	0.28
Agents Held	1.35	1.63	2.05	1.14	1.45	1.78
Agents In-call	0.26	0.31	0.36	0.24	0.29	0.34
Msec per BIO Request	46.87	51.91	55.69	47.04	50.72	55.79
Blocking Factor (Blocks/BIO)	1.36	1.36	1.34	1.34	1.35	1.33
Chaining Factor (Blocks/IO)	1.32	1.31	1.30	1.32	1.33	1.32

B.1.6 EDF and SFS Comparisons

SECTION	4.3.2.1	4.3.2.2	4.3.2.2	4.3.2.2
RELEASE	ESA 1.0	ESA 1.0	ESA 1.0	ESA 1.0
RUN ID	EC7031	EC7029	EC7029	EC7029
PROCESSOR	4381-13	4381-13	4381-13	4381-13
STORAGE	16M	16M	16M	16M
WORKLOAD	MAX% SFS	35% SFS	35% SFS	35% SFS
ACTIVE USERS	63	69	77	85
NORMALIZED BY COMMAND				
Checkpoints Taken	0.001	0.001	0.001	0.001
Checkpoint Time (msec)	1.557	2.099	1.520	3.014
Close File Requests	0.749	0.442	0.432	0.442
Connect Requests	0.006	0.004	0.004	0.003
Delete File Requests	0.104	0.107	0.105	0.106
File Copy Requests	0.003	0.000	0.000	0.000
Lock Requests	0.027	0.027	0.027	0.027
Open File New Requests	0.000	0.002	0.002	0.002
Open File Read Requests	0.566	0.254	0.246	0.255
Open File Replace Requests	0.159	0.161	0.161	0.160
Open File Write Requests	0.025	0.025	0.025	0.025
Query File Pool Requests	0.000	0.000	0.000	0.000
Query User Space Requests	0.023	0.023	0.023	0.023
Read File Requests	0.387	0.255	0.254	0.256
Refresh Directory Requests	0.029	0.014	0.015	0.014
Rename Requests	0.005	0.005	0.005	0.005
Unlock Requests	0.026	0.027	0.027	0.027
Write File Requests	0.145	0.141	0.141	0.142
Total File Pool Requests	2.254	1.486	1.466	1.488
File Pool Request Service Time (msec)	146.803	120.856	145.250	238.692
Local File Pool Requests	2.254	1.486	1.466	1.488
Begin LUWs	0.844	0.555	0.547	0.553
Agent Holding Time (msec)	410.814	366.066	472.838	777.715
SAC Calls	9.729	6.912	6.837	6.846
Catalog Lock Conflicts	0.002	0.002	0.002	0.004
Total Lock Conflicts	0.002	0.002	0.002	0.004
Lock Wait Time (msec)	0.103	0.135	0.057	0.275
File Blocks Read	1.718	1.053	1.048	1.059
File Blocks Written	0.652	0.644	0.639	0.647
Catalog Blocks Read	1.177	0.668	0.710	0.777
Catalog Blocks Written	0.320	0.310	0.324	0.347
Ctl Minidisk Blocks Written	0.021	0.023	0.018	0.021
Log Blocks Written	0.594	0.607	0.599	0.588
Total DASD Block Transfers	4.482	3.306	3.338	3.439
BIO Requests to Read File Blocks	0.992	0.541	0.540	0.548
BIO Requests to Write File Blocks	0.311	0.309	0.309	0.313
BIO Requests to Read Catalog Blocks	1.177	0.668	0.710	0.777
BIO Requests to Write Catalog Blocks	0.320	0.310	0.324	0.347
BIO Requests to Write Ctl Minidisk Block	0.004	0.005	0.004	0.004
BIO Requests to Write Log Blocks	0.594	0.607	0.599	0.588
Total BIO Requests	3.398	2.440	2.486	2.578
Total BIO Request Time (msec)	98.137	80.493	105.878	162.928
I/O Requests to Read File Blocks	0.993	0.542	0.541	0.550
I/O Requests to Write File Blocks	0.311	0.309	0.309	0.313
I/O Requests to Read Catalog Blocks	1.177	0.668	0.710	0.777
I/O Requests to Write Catalog Blocks	0.320	0.310	0.324	0.347
I/O Requests to Write Ctl Minidisk Block	0.006	0.007	0.006	0.006
I/O Requests to Write Log Blocks	0.594	0.607	0.599	0.588
Total I/O Requests	3.402	2.443	2.489	2.581
DERIVED RELATIONSHIPS				
Deadlocks	0	0	0	0
Rollbacks Due to Deadlock	0	0	0	0
LUW Rollbacks	0	0	0	0
Checkpoints Taken	5	6	5	6
Checkpoint Duration (sec)	2.17	2.62	2.48	4.19
Seconds Between Checkpoints	241.00	201.33	241.40	201.67
Checkpoint Utilization	0.90	1.30	1.03	2.08
Agents Held	2.38	2.27	3.19	5.36
Agents In-call	0.85	0.75	0.98	1.64
Msec per BIO Request	28.88	32.99	42.60	63.20
Blocking Factor (Blocks/BIO)	1.32	1.36	1.34	1.33
Chaining Factor (Blocks/IO)	1.32	1.35	1.34	1.33

B.1.7 4381-13 / 32M

SECTION	4.7.2.1	4.7.2.1	4.7.2.1
RELEASE	ESA 1.0	ESA 1.0	ESA 1.0
RUN ID	EC7605	EC7605	EC7605
PROCESSOR	4381-13	4381-13	4381-13
STORAGE	32M	32M	32M
WORKLOAD	35% SFS	35% SFS	35% SFS
ACTIVE USERS	59	67	75
NORMALIZED BY COMMAND			
Checkpoints Taken	0.001	0.001	0.001
Checkpoint Time (msec)	0.837	0.817	0.990
Close File Requests	0.426	0.441	0.435
Connect Requests	0.004	0.004	0.004
Delete File Requests	0.104	0.106	0.104
Lock Requests	0.027	0.027	0.027
Open File New Requests	0.002	0.002	0.002
Open File Read Requests	0.240	0.254	0.248
Open File Replace Requests	0.160	0.162	0.160
Open File Write Requests	0.025	0.025	0.025
Query File Pool Requests	0.000	0.000	0.000
Query User Space Requests	0.023	0.024	0.023
Read File Requests	0.252	0.253	0.255
Refresh Directory Requests	0.014	0.015	0.015
Rename Requests	0.005	0.005	0.005
Unlock Requests	0.027	0.027	0.027
Write File Requests	0.142	0.139	0.140
Total File Pool Requests	1.451	1.484	1.470
File Pool Request Service Time (msec)	78.013	82.034	92.325
Local File Pool Requests	1.451	1.484	1.470
Begin LUWs	0.538	0.557	0.544
Agent Holding Time (msec)	250.801	262.722	304.727
SAC Calls	6.764	6.910	6.774
Catalog Lock Conflicts	0.001	0.003	0.002
Total Lock Conflicts	0.001	0.003	0.002
Lock Wait Time (msec)	0.014	0.052	0.038
File Blocks Read	1.034	1.052	1.051
File Blocks Written	0.645	0.638	0.641
Catalog Blocks Read	0.607	0.637	0.673
Catalog Blocks Written	0.279	0.298	0.307
Ctl Minidisk Blocks Written	0.021	0.019	0.022
Log Blocks Written	0.606	0.609	0.596
Total DASD Block Transfers	3.191	3.254	3.289
BIO Requests to Read File Blocks	0.527	0.543	0.541
BIO Requests to Write File Blocks	0.309	0.306	0.308
BIO Requests to Read Catalog Blocks	0.607	0.637	0.673
BIO Requests to Write Catalog Blocks	0.279	0.298	0.307
BIO Requests to Write Ctl Minidisk Block	0.004	0.004	0.004
BIO Requests to Write Log Blocks	0.606	0.609	0.596
Total BIO Requests	2.331	2.397	2.429
Total BIO Request Time (msec)	51.992	54.203	62.479
I/O Requests to Read File Blocks	0.528	0.544	0.543
I/O Requests to Write File Blocks	0.309	0.306	0.308
I/O Requests to Read Catalog Blocks	0.607	0.637	0.673
I/O Requests to Write Catalog Blocks	0.279	0.298	0.307
I/O Requests to Write Ctl Minidisk Block	0.007	0.006	0.006
I/O Requests to Write Log Blocks	0.606	0.609	0.596
Total I/O Requests	2.334	2.401	2.433
DERIVED RELATIONSHIPS			
Deadlocks	0	0	0
Rollbacks Due to Deadlock	0	0	0
LUW Rollbacks	0	0	0
Checkpoints Taken	5	5	6
Checkpoint Duration (sec)	1.13	1.24	1.38
Seconds Between Checkpoints	240.60	240.80	200.83
Checkpoint Utilization	0.47	0.51	0.69
Agents Held	1.40	1.65	2.12
Agents In-call	0.44	0.51	0.64
MSec per BIO Request	22.30	22.61	25.73
Blocking Factor (Blocks/BIO)	1.37	1.36	1.35
Chaining Factor (Blocks/IO)	1.37	1.36	1.35

Appendix C. Workloads

C.1 CMS Interactive (FS7B)

C.1.1 History of the CMS Interactive Workload

For several years, the performance departments have relied on a general interactive fullscreen workload for multi-user measurements of VM operating systems. This workload consists of a number of CMS users performing various tasks (editing, compiling, program execution, etc.). The CMS users are simulated by the Full Screen Internal Driver (FSID) tool which simulates graphic terminal inputs (3270) as if they were entered by an actual user. FSID provides a mechanism to automatically execute command scripts in virtual machines without the need for physical terminals or human operators.

The workload, known as FSMIX3B, consisted of 18 scripts, each performing a specific function, that were executed by CMS users running on a VM operating system. The FSMIX3B workload was initially run during performance testing of VM/SP Release 3 and, with minor modifications, against releases 4, 5 and 6. It was used primarily for regression measurements, comparing the performance of a new VM release against the performance of the previous release. For a valid regression comparison, it is necessary to run identical workloads on identical hardware configurations.

In VM/SP Release 6, a major new function, Shared File System (SFS), was added. Since this function changes the way data files are handled within VM, it was not possible to directly compare Release 6 (with SFS) to Release 5 (with EDF), using the same workload. A version of the workload (FSMIX6D) was built which substituted shared file activity for the corresponding minidisk activity in previous versions. This workload was then used to evaluate the performance of SFS versus the minidisk system.

Since the Shared File System has become a major part of VM/SP, it was decided to investigate the possibility of upgrading FSMIX3B to include some level of SFS activity as a part of the base workload, while maintaining the ability to measure against minidisks. From this need came the FS7B workload.

C.1.2 Reasons for Change

A series of meetings were held to solicit suggestions for improvements to the basic workload. Among the suggestions, several were deemed to be major. They include:

1. Generally update existing workload scripts - Incorporate changes that reflect functions added during the last several releases of VM/SP as well as functions that the existing workload lacked. Examples include:

   Shared File System activity

   Full Screen CMS activity

   Help processing

   CMS Macros in Assembler scripts

   LOGON/LOGOFF activity

   Exec processing

2. Create an incentive for improving Shared File System performance - A major responsibility of the performance departments is to recommend changes that will result in performance improvements. We have found that it is easier to prove the value of a recommended change if the change impacts the performance of the basic regression workload.

3. Keep an incentive for maintaining EDF performance - Although the Shared File System is a growing part of VM, minidisk I/O will remain a major part of system performance for the foreseeable future. It is still important to improve its performance.

4. Increase the load (CPU, Storage, etc.) - As the size of supported processors increases, it was felt that our original workload should be heavier to be representative of the typical VM environment. Increasing the overall load provides an alternative to simply increasing the number of users when scaling up to larger systems.

5. Increase read to write ratio - To have a more realistic customer environment, it was necessary to increase the number of reads over writes. The FSMIX3B workload was too low at 1.5 to 1. A typical environment is closer to 4 to 1, which the FS7B workload provides.

6. Create one unique workload for multiple purposes - In order to meet the productivity goals, it is important to minimize the number of machine runs that are needed to obtain necessary information. By changing the variables for the search order in one file, the workload can be transformed from all minidisk to all SFS or in between.

7. Create one unique workload for VM - As VM grew, it became necessary to create one unique workload that runs across the VM line (SP, HPO, XA, ESA). The data can be compared from the low end (9370) to the high end (3090).

C.1.3 What is FS7B?

It was decided that the new workload should be based on the FSMIX5H/ FSMIX6D version of the basic workload. These are the minidisk and shared file versions that were used to evaluate VM/SP Release 6 against VM/SP Release 5. The FS7B workload has been designed so that it can have all minidisk activity (FS7B0), all SFS activity except for the S and Y disk (FS7BMAX), or a mixture of both. The target for Shared File System activity was set at 35% of user file I/Os (FS7B35). This would be achieved by placing all end user files into SFS. A windowing script was added to exercise the CMS fullscreen function, added in VM/SP Release 5. This script also exercises the IPL CMS command and the LOGON/LOGOFF activity. Another new script includes HELP processing.

It was decided that four of the original scripts, two IBM BASIC and two APL, could be dropped because they had contributed little to the finding of problems in past releases and were probably not representative of a typical CMS user profile. This also allowed us to keep the number of scripts to a manageable level for run repeatability (< 20). Other changes made that were general cleanup types were:

Eliminate fixed delays in XEDIT scripts

Break up large inputs in XEDIT scripts

Reduce use of ERASE and COPYFILE

Increase usage of REXX

Include H-Assembler as well as F-Assembler

Include CMS Macros in Assembler scripts

Update Program Product levels to the SP6 System Offering Level

Increase the Read/Write ratio

Include LOGON/LOGOFF activity

Expand search order

Each user's search order was expanded to include nine accessed disks or directories. It was felt that the search order in the old workload version, consisting of five minidisks, was not representative of a typical CMS user.

The FS7B workload can be used for many types of runs just by changing the search order. It can be an all minidisk workload by accessing all minidisks. It can be an all SFS workload by accessing all directories. It can also be a combination of both. A file called DSKORDER EXEC has been added to the user's Y-disk (MAINT'S 19E) with a filemode of Y2. This file is executed by the PROFILE EXEC on the user's A-disk. At the beginning of the DSKORDER EXEC, there are seven variables, named ADISK, BDISK, CDISK, DDISK, EDISK, FDISK, and GDISK. By changing the variable to either 'MINI' for minidisk or 'SFS' for SFS directories, many different combinations of the workload can be executed. Below are how the fields look for the three basic workloads:

```
  FS7B0            FS7B35           FS7BMAX
  -----            ------           -------
ADISK='MINI'     ADISK='SFS'      ADISK='SFS'
BDISK='MINI'     BDISK='SFS'      BDISK='SFS'
CDISK='MINI'     CDISK='MINI'     CDISK='SFS'
DDISK='MINI'     DDISK='SFS'      DDISK='SFS'
EDISK='MINI'     EDISK='MINI'     EDISK='SFS'
FDISK='MINI'     FDISK='MINI'     FDISK='SFS'
GDISK='MINI'     GDISK='MINI'     GDISK='SFS'
```

To handle the logon of users after logging off, a special server, named LOGONSRV, has been included. During the run, this user will be run NEWLOGON EXEC while disconnected. This exec uses the WAKEUP module. The only activity LOGONSRV will have is if it receives an SMSG (special message) from a user that logs off or if it is awakened by the WAKEUP module. When a user logs off, it sends a SMSG containing its GRAF-ID and USERID. After one minute, this user will be logged-on again to its original GRAF-ID.

Note: Only the measurements that run with FSID use the above method to handle users that logon/logoff. TPNS measurements handle the logon/logoff in one of the workload scripts (WND517L).

C.1.4 Workload Descriptions

C.1.4.1 FS7B Variations

By changing the search order of the FS7B workload, different combinations can be run to test many environments ranging from all minidisk activity to all SFS activity or anywhere in between. Below are the six different FS7B workloads that were used for regression testing.

FS7B0 Workload: All filemodes are accessed as minidisk. There is no SFS usage. Local users are simulated with FSID.

FS7B0R Workload: All filemodes are accessed as minidisk. There is no SFS usage. Remote users are simulated with TPNS.

FS7B35 Workload: SFS directories are accessed as filemodes A, B, and D. All other filemodes are accessed as minidisk. Approximately 35% of all minidisk I/Os are eliminated as the activity that caused them is assumed by the Shared File System. Local users are simulated with FSID.

FS7B35R Workload: SFS directories are accessed as filemodes A, B, and D. All other filemodes are accessed as minidisk. Approximately 35% of all minidisk I/Os are eliminated as the activity that caused them is assumed by the Shared File System. Remote users are simulated with TPNS.

FS7BMAX Workload: All filemodes are accessed as SFS directories except S and Y. Approximately 48% of all minidisk I/Os are eliminated as the activity that caused them is assumed by the Shared File System. Local users are simulated with FSID.

FS7BMAXR Workload: All filemodes are accessed as SFS directories except S and Y. Approximately 48% of all minidisk I/Os are eliminated as the activity that caused them is assumed by the Shared File System. Remote users are simulated with TPNS.

C.1.4.2 FS7B Program Products

The following program products are used by the FS7B Workload.

```
       COBOL VS 2 - Version 1 Release 3.0
       DCF        - Version 1 Release 3.2 (Shared Segments)
       FORTRAN VS - Version 2 Release 3.0 (Shared Segments)
       HASM       - Version 2 Release 1.0
>      PL/I       - Version 2 Release 1.0
```

C.1.4.3 FS7B Virtual Machine Configuration

This section describes the virtual machine configuration used by the FS7B Workload.

- Below is the search order used for the FS7B workload:

```
A-disk - R/W and contained 100 files (Root directory for SFS)
B-disk - R/W and contained 0 files (Sub-directory for SFS)
C-disk - R/O and contained 500 files (Operator's directory for SFS)
D-disk - R/W and contained 500 files (Operator's directory for SFS)
E-disk - R/O and contained 500 files (Operator's directory for SFS)
F-disk - R/O and contained 500 files (Operator's directory for SFS)
G-disk - R/O and contained 500 files (Operator's directory for SFS)
S-disk - R/O and contained approx. 240 files
Y-disk - R/O and contained approx. 400 files (PID-level msmts.)
Y-disk - R/O and contained approx. 750 files (early msmts.)
```

- All files on the C-disk had the FSTs saved in a shared segment for the minidisk only measurements.
- HELP disk had the FSTs saved in a shared segment.
- CMSINST and CMSVMLIB shared segments were used.
- CMSFILES shared segment was used when SFS was used.
- SFS Granted Authority to PUBLIC Read only for all of Operator's directories except the D-disk which was Granted Write Authority.

C.1.4.4 FS7B Script Summary

The FS7B Workload consists of seventeen scripts plus an initialization script. This script (INIT7 for FSID or LOGESA for TPNS) is executed once by each user at LOGON time to setup the needed file structure and CMS configuration. The scripts are:

Script Name	% Used	Script Description
INIT7	0%	Initialization (FSID)
LOGESA	0%	Initialization (TPNS)
ASM617	5%	BAL Assemble (HASM) and Execution
ASM627	5%	BAL Assemble and Execution
XED117	5%	EDIT of a VSBASIC Program
XED127	10%	EDIT of a VSBASIC Program
XED137	10%	EDIT of a COBOL Program
XED147	10%	EDIT of a COBOL Program
COB217	5%	COBOL Compile
COB417	5%	Execute a COBOL Program
FOR217	5%	VSFORTRAN Compile
FOR417	5%	FORTRAN Execution
DCF517	5%	Edit and Script a File
PRD517	5%	Productivity Aids Session
PLI317	5%	PL/I Optimizer Session
PLI717	5%	PL/I Optimizer Session
WND517	8%	Window Exploitation with IPL CMS
WND517L	2%	Window Exploitation with LOGON/LOGOFF
HLP517	5%	HELP Exploitation

C.1.4.5 FS7B Script Modifications

During the course of regressions tests from the base to VM/ESA, new functions were added or modified in VM/ESA which resulted in the modification of some of the workload scripts. The scripts still perform the identical functions but are just implemented in a different manner. Below are the scripts which were affected by the changes implemented by VM/ESA.

- ASM617 and ASM627

  The base runs do a global of 2 maclibs (DMSSP and CMSLIB) before assembling the source. In VM/ESA, only 1 maclib (DMSGPI) will be globaled before assembling the source.

  For the HPO 5.0 measurements, the A100A and A100C Assemble and Text files did not use the CMSSTOR macro since it did not exist. The DMKFREE macro was substituted for it.

- WND517 and WND517L

  The base runs use the original two word fullscreen commands (POP WINDOW MESSAGE, DROP WINDOW MESSAGE, etc.). The base runs also do a HELP WINDOW which gives a menu of commands to choose from. In the base runs, PF1 is used to choose CHANGE WINDOW SIZE. In VM/ESA, new two word fullscreen commands were implemented (WINDOW POP MESSAGE, WINDOW DROP MESSAGE, etc.). Also, the HELP WINDOW does not give a menu anymore. Therefore, instead of the PF1 command, the command HELP WINDOW SIZE is issued at the command line.

C.1.4.6 FS7B Script Descriptions

The following is a summary of each script used for the FS7B workload.

INIT7: Initialization Script (FSID)

General Description
Every user executes this script first to set up the virtual machine.

Summary of 2 Script Commands
Execute DELEX Exec to clean A-disk.
Execute PROFILE Exec to set correct search order, set acnt off, set printer class d, and set terminal linend off.

LOGESA: Initialization Script (TPNS)

General Description
Every user executes this script first to set up the virtual machine.

Summary of 7 Script Commands
Set autoread on.
Execute CHKFTMA exec to check the format of the A-disk.
Execute CHKFTMB exec to check the format of the B-disk.
Execute DELEX Exec to clean A-disk.
Access 191 as A-disk.
Execute PROFILE Exec to set correct search order, set acnt off, set printer class d, and set terminal linend off.
Set remote on.

ASM617: BAL Assemble (HASM) and Execution

General Description
This is an assembly, using HASM, and execution of a 125 statement program with 675 comment lines.

Summary of 24 Script Commands
Query reader and printer.
Spool printer class D.
Xedit A100A ASSEMBLE and qquit.
Global appropriate maclibs.
Listfile A100C ASSEMBLE.
Assemble the source using HASM (NOLIST option).
Erase the text deck.
Repeat the above 1 more time except for xedit.
GLOBAL maclib reset.
Load the text file (NOMAP option).
Generate a module (NOMAP option).
Execute the module.
Load the text file (NOMAP option).
Execute the module 2 more times.
Execute DELEX Exec to clean A-disk.

ASM627: BAL Assemble and Execution

General Description

This is an assembly, using the F-Assembler, and execution of a 125 statement program with 675 comment lines.

Summary of 21 Script Commands

Query reader and printer.
Spool printer class D.
Global appropriate maclibs.
Listfile A100C ASSEMBLE.
Xedit A100C ASSEMBLE and qquit.
Assemble the source (NOLIST option).
Erase the text deck.
Repeat the above 1 more time except for xedit.
GLOBAL maclib reset.
Load the text file (NOMAP option).
Generate a module (NOMAP option).
Execute the module.
Load the text file (NOMAP option).
Execute the module.
Repeat the last 2 statements 1 more time.
Execute DELEX Exec to clean A-disk.
Query disk, users, and time.

XED117: Edit of a VSBASIC Program

General Description

The script uses XEDIT in full screen mode to update an existing VSBASIC program. The program consists of 69 statements.

Summary of 32 Script Commands

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

XED127: Edit of a VSBASIC Program

General Description

This uses XEDIT in full screen mode to edit a VSBASIC program.

Summary of 30 Script Commands

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 (QQUIT).
Quit FILELIST.
Repeat all of the above statements, changing 9 lines instead of 6 and without issuing the top and bottom commands.

XED137: Edit of a COBOL Program

General Description

This is an edit of a 387 statement COBOL program using XEDIT in full screen mode.

Summary of 30 Script Commands

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

XED147: Edit of a COBOL Program

General Description

This is an edit of a 387 statement COBOL program using XEDIT in full screen mode.

Summary of 31 Script Commands

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

COB217: COBOL Compile

General Description

This script compiles a 395 statement COBOL program.

Summary of 29 Script Commands

Set ready message short.
Link and access a disk.
Query link and disk.
LISTFILE the program.
Invoke the COBOL compiler.
Erase the compiler output.
Release and detach the linked disk.
Set ready message long.
Set message off.
Query set.
Set message on.
Set ready message short.
Link and access a disk.
LISTFILE the program.
Invoke the COBOL compiler.
Erase the compiler output.
Release and detach the linked disk.
Query term and rdymsg.
Set ready message long.
Set message off.
Query set.
Set message on.
Purge printer.

COB417: Execute a COBOL Program

General Description

This script executes a COBOL program under CMS. The program contains 410 source statements.

Summary of 28 Script Commands

Define T-disk space for 2 disks using an exec.
Query dasd and format both T-disks.
Establish FILEDEFs for input and output files (4).
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 T-disks.
Define one T-disk as another.
Detach the T-disks.
Reset the GLOBAL and clear the FILEDEFs.

FOR217: VS FORTRAN Compile

General Description

This is a compile of 6 VS Fortran programs.

Summary of 23 Script Commands

Nucxdrop Namefind w/ NUKENAME Exec
Query and purge the reader.
Compile UOFIH2.
Issue indicate commands.
Compile UOFIH4.
Issue indicate commands.
Compile UOFIH1.
Issue indicate commands.
Repeat the above 6 statements.
Execute DELEX Exec to clean A-disk.
Purge the printer.

FOR417: FORTRAN Execution

General Description

This is an execution of 3 Fortran programs.

Summary of 27 Script Commands

Spool printer class D.
GLOBAL appropriate text libraries.
Issue two FILEDEFs for output.
Load and start UOFIH2 (NOMAP option).
Rename output file and purge printer.
Repeat above 5 statements for UOFIH1 and UOFIH4, except
erase the output file for UOFIH1 and don't issue spool printer.
List and erase output files.
Reset GLOBAL and clear FILEDEFs.

DCF517: Edit and Script a File

General Description

This script uses XEDIT in full screen mode to enter a document, then uses DCF to format and display it on the terminal.

Summary of 31 Script Commands

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

PRD517: Productivity Aids Session

General Description

A session that makes use of the following : REXX, NAMES, SENDFILE, PEEK, RECEIVE, DISCARD and RDRLIST.

Summary of 22 Script Commands

Execute MYID6 Exec.
Issue NAMES command and add operator.
Locate a user in names file and quit.
Issue the SENDFILE command.
Send a file to * (yourself).
Repeat the above 2 statements twice but send the file to ME (yourself).
Issue RDRLIST command, PEEK and DISCARD a file.
Refresh RDRLIST screen, Receive a file on B-disk, and quit.
Transfer all Reader files to Punch.
Purge Reader and Punch.
Execute a REXX exec that generates 175 random numbers.
Execute a REXX exec that reads multiple files of various sizes from
both the A-disk and C-disk.
Erase MYID6 EXEC off B-disk.
Execute DELEX Exec to clean A-disk.

PLI317: PL/I Optimizer Session

General Description

XEDIT in full screen mode and compile a PL/I Optimizer program with 101 statements.

Summary of 28 Script Commands

Do a GLOBAL txtlib.
Perform a FILELIST.
XEDIT the program.
Execute 15 XEDIT subcommands.
File the results on A-disk with a new name.
Quit filelist.
Enter two FILEDEFs for compile.
Compile it using PLIOPT.
Erase the PL/I program.
Reset the GLOBAL and clear the FILEDEFs.
Query virtual devices.
Tell * (yourself) one pass of script executed.

PLI717: PL/I Optimizer Session

General Description

XEDIT in full screen mode, compile and execute a PL/I Optimizer program of 47 statements.

Summary of 27 Script Commands

Copy and rename the PL/I program and data file from C-disk.
XEDIT data file and QQUIT.
Xedit IAWN1 File.
Issue right 20, left 20, and set verify on.
Change two lines.
Change filename to IAWN and file the result.
Compile using PLIOPT.
Set two FILEDEFs and query the settings.
GLOBAL for PL/I transient library.
Load the program (NOMAP option).
Start the program.
Type 8 lines of one data file.
Execute DELEX Exec to clean A-disk.
Erase extra files on B-disk.
Reset the GLOBAL and clear the FILEDEFs.
Tell * (yourself) one pass of script executed.

WND517: Window Expliotation

General Description

Exploits window commands with fullscreen on and IPL CMS.

Summary of 28 FSID or 30 TPNS Script Commands

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 (only for TPNS).
Set remote on (only for TPNS).

WND517L: Window Exploitation

General Description

Exploits window commands with fullscreen on and LOGOFF.

Summary of 28 FSID or 31 TPNS Script Commands

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.
Execute PROFLOGF Exec to send SMSG to LOGONSRV and LOGOFF (only for FSID).
Logoff user and wait 60 seconds (only for TPNS).
Logon user back to original GRAF-ID (only for TPNS).
Set autoread on (only for TPNS).
Set remote on (only for TPNS).

HLP517: Help Exploitation

General Description

Exploits HELP and other Misc. commands.

Summary of 28 Script 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.
Execute DELEX Exec to clean A-disk.

C.2 Connectivity

C.2.1 Connectivity Driver Tool

The Connectivity Driver consists of two parts: 1) the User part (Application/Requestor) and 2) the Resource part (Server). Each part runs in its own virtual machine. The basic scenario is as follows:

- Initialization Process

    1. The Resource program is started. It identifies itself to CP as a resource and waits for APPC/VM interrupts.

    2. The User program is invoked, issuing a CONNECT to the Resource. The number of bytes to transfer (n) is specified as a parameter when the user program is invoked.

    3. The Resource accepts the connection; an APPC/VM connection (path) has been created.

    4. The Resource issues APPC/VM RECEIVE for that User. The Resource is ready to receive requests.

- Send/Receive Process (process that is measured)

    1. The User application prompts the user for the number of iterations (i) to transfer n bytes.

    2. User application issues an APPC/VM SENDDATA with RECIEVE=YES and an ANSBUF (area to put Resource's reply). Approximately 25 bytes (making up the request) are sent to the Resource.

    3. The Resource's RECEIVE (or SENDDATA RECEIVE=YES) completes; it now has a request from a User for n bytes.

    4. The Resource issues an APPC/VM SENDDATA RECEIVE=YES of n bytes (plus a few control bytes) to the User.

    5. The User's RECEIVE completes for the n bytes.

    6. Steps 2 thru 5 are performed i times

It is important to note the data transfers described above are from memory to memory. Because of this, there is no disk I/O once the program is read in.

The Connectivity Driver uses the CMSIUCV interface to set up APPC connections. We use a CMS Communications Directory so that the driver doesn't have to change when running through AVS.

C.2.2 Workload Descriptions

- CONN1 Workload

    The CONN1 workload is the single thread connectivity workload. It uses the Connectivity Driver Tool to exercise APPC/VM data transfers. See "Connectivity Driver Tool" for an explanation of the Connectivity Driver Tool. The Connectivity Driver is invoked for multiple iterations of data transfers. The number of iterations is chosen such that the overhead of command initialization is negligible when resources per iteration are computed.

- Multi-user Workload

    The Multi-user workload uses the Connectivity Driver Tool to exercise APPC/VM data transfers. See "Connectivity Driver Tool" for an explanation of the Connectivity Driver Tool.

There can be multiple User virtual machines connecting to a single Resource virtual machine. There can be any number of Resource virtual machines. Most (but not all) of the runs documented in this report used one Resource machine for every real CPU in the processor complex.

The TSAF and APPC/VM measurements each use the same set of FSID scripts. The AVS measurements change the initialization script so that the User loads a CMS Communications Directory file.

C.2.3 Script Descriptions

All the scripts are trivial variations of the following two:

- Connect

  Every User executes this script to set up his/her virtual machine for 512-byte transfer. This script is not measured.

  Summary of Script Commands

  Set accounting off.
  Spool printer class d.
  Term linend off.
  If this is an AVS script, load a CMS Communications Directory so
  that the resource name is resolved correctly.
  Establish APPC/VM path to Resource for a given size of data transfer

- Transfer

  Every APPC/VM or TSAF user executes this script to transfer data.

  Summary of Script Commands

  Get given size of data i times.

C.3 CRRTOOL

The Coordinated Resource Recovery Tool (CRRTOOL) is an Internal Use Only tool which measures CRR performance. The tool came into existence since no CMS applications were available and no test applications existed which would test CRR. Also, there were no CMS commands which exercised the two-phase commit.

The tool gives the user the ability to test many scenarios among multiple resources. For example, a user can perform any combination of updates or reads to a maximum number of three filepools. All updates are either committed or rolled back.

The tool consists of two Assembler programs, CRRTOOLU which runs in the user machine and CRRTOOLR which runs in the server machine. CRRTOOLU can be run by one or more users on the test system. CRRTOOLU has two levels of coordination: one-phase commit and two-phase commit.

CRRTOOLU uses protected conversation (APPC/VM) to communicate with the CRRTOOLR service machine. CRRTOOLR acts as the partner in the protected conversation for the CRRTOOLU application. Below is a sample CRR configuration.

CRRTOOL Example Configuration

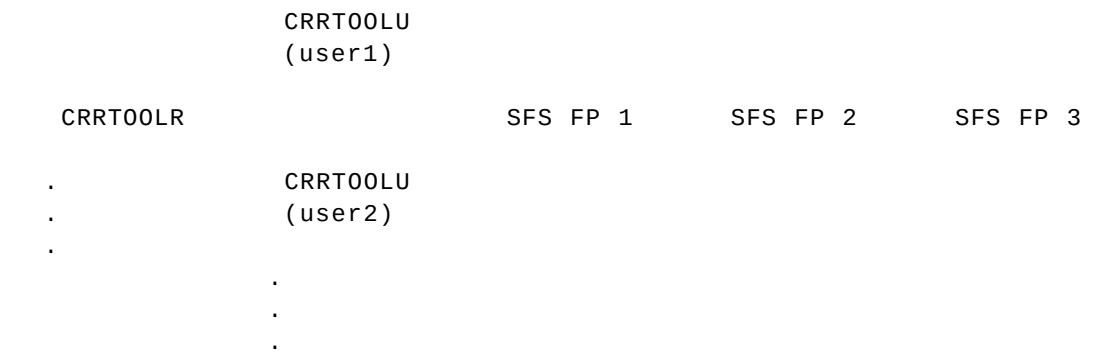

An example CRRTOOL configuration. Note there could be any number of CRRTOOLU virtual machines and multiple CRRTOOLR virtual machines.

The basic scenario for the CRRTOOL is as follows:

Initialization Process

1. The inputs to CRRTOOLU are parsed.
2. If message is active, then write out the start up messages and user parameters. Message inactive is preferred when performance testing to eliminate overhead.
3. Check if protected conversation is active or not. If protected conversation is active then extablish connections between CRRTOOLU and CRRTOOLR.
4. Initialize filepool 1, 2, and 3 to read, update or not active.

Measured Process

1. Prompt the user for the number (n) of syncpoints to process per transaction. If a zero is entered, the application ends.
2. If protected conversation is active then send data from one application to the other.
3. If filepool 1, 2, or 3 is read then call the CSL function DMSEXIFI (checks for the existence of the file). If filepool 1, 2, or 3 is update then call the CSL function DMSRENAM (renames the file).
4. Commit the work by issuing the CSL function DMSCOMM (commits filepool changes).
5. Steps 1 through 4 are repeated n times.

For more information on the CRRTOOL, please refer to the ***CRRTOOL User Guide and Reference.***

C.4 MVS Guest (CB84)

C.4.1 MVS Guest (CB84) Workload Description

CB84 (Commercial Batch 1984) is a jobstream intended to represent an MVS/XA commercial batch workload. It is made up of a variety of customer programs, utilities, and synthetic jobs. One copy of the CB84 workload contains a total of 130 batch jobs that contain 610 job steps and use 1,021 permanent data sets. Fifty one of the jobs are unique, while the remaining seventy nine are replications. All of the job steps execute programs except for two steps that execute instream procs. The following is a breakdown of the jobs contained in one copy of the CB84 workload:

38 COBOL Go Jobs
15 COBOL Compile and LINKEDIT Jobs
15 IEBJENNER Jobs
15 BAL Assemble and LINKEDIT Jobs
14 PL/I Go Jobs
14 Synthetic Jobs
10 PL/I Compile and Go Jobs
7 IEBCOPY and COMPRESS Jobs
1 BAL Go Job
1 COBOL Compile and Go Job

The executed programs include inventory, banking, payroll and table update applications, as well as synthetic jobs that do fixed point arithmetic, GETMAINs, FREEMAINs and private storage area references designed to represent those observed in customer workloads. Many of the jobs do heavy I/O and make extensive use of multiple data sets and libraries. Paging activity is minimal to non-existent.

C.4.2 Measurement Methodology

For VM/ESA 1.0 measurements, preliminary CB84 runs were required to ′prime′ VLF with the appropriate modules. Data and tuning information from these initial runs were not valid as measurement data.

The measurements began by tuning MVS (native) to determine reasonable values for the number of initiators to start and the number of copies of jobs to run to keep the system busy for at least 10 minutes. The number of initiators depended on the workload and on the I/O configuration. It was actually set, by experimentation, to the number that resulted in processor utilization of at least 100% during the steady-state portion of the workload execution. This information was available in the ″CCVUTILP AVERAGE″ field of the RMF Trace Activity report. The overall utilization had to be at least 80%. The batch jobstream was released when the RMF ZZ ACTIVE message appeared by using a PF key rather than the $VS command. The RMF data was inspected to ensure I/O balancing and to compute ITR and ETR values to determine the maximum thruput. The system was then loaded with jobs while the queues were held to allow all preliminary work to complete before starting the measurement. The next step was to simultaneously release the queues and start RMF and other measurement tools, e.g., MONITOR, RTM and an internal Counters program.

The measurements were run twice under the same conditions to validate the results and to show that they were repeatable.

C.4.3 Criteria for Valid Measurements

The following is a list of the items that must be checked in order to validate a CB84 measurement.

- High utilization greater than or equal to 80%

  RMF trace activity report had to show at least 100% in steady state.

- Less than 5 I/O errors over a 10-minute time span
- No abends
- No permanent I/O errors
- No missing I/O interrupts
- Items to look at when comparing runs:

  Internal throughput (jobs/CPU sec)
  Total instructions
  Supervisor instructions
  Problem program instructions
  Total instructions/CACHE Miss
  Total instructions/job
  Supervisor instructions/job
  Problem program instructions/job
  Pages/job (total, demand, swap)
  DASD IOs/job
  TCB CPU SUs/job
  No greater than 1% variation in problem program instructions/job

- Clean erep

Appendix D. Configuration Details

D.1 DASD Configurations

D.1.1 Introduction

There are various DASD volumes associated with each processor in this report. Each DASD is laid out for good performance and I/O load balancing. A portion or all of each DASD may be used for the various workloads. For instance, if the FS7B0 workload is running on the 4381-13; only PAGE, USERDISKS, SPOOL, and TDISK space will be used on MDSK01. For the SFS measurements, such as the FS7B35 workload, then the same areas mentioned above along with SFS Control, SFS Storage Group 1, and SFS Storage Group 2 will be used.

D.1.2 Key Terms and Abbreviations

The following list below contains the abbreviations and descriptions of each key item on every DASD volume used for performance measurements.

```
CH          - Channel.
CPRES       - The resident volume of the CP nucleus.
CTC         - Channel-to-Channel Adapter.
HELP        - HELP disk.
MVS/IMS     - Multiple Virtual Storage / Information Management System.
MVS/IPL     - Multiple Virtual Storage / Initial Program Load.
OFFICE      - Office Vision.
OPDISKS     - Operator's 295, 296, 395 disks used for FS7B.
PAGE        - CP area used for page space.
SPOOL       - CP area used for spool space.
SWAP HPO    - Swapping packs used for HPO.
S-DISK      - CMS system disk.
TDISK       - CP area used for temporary disks.
TYPE        - Type of DASD (9332, 9335, 3380).
USERDISKS   - 191 and 111 disks for each of the users.
VOLID       - Name of the DASD pack.
xxx CONTROL - Control minidisk used by SFS or CRR servers.
xxx CRR1    - CRR log disk used by the CRR server.
xxx CRR2    - CRR log disk used by the CRR server.
xxx LOG1    - Log minidisk used by SFS or CRR servers.
xxx LOG2    - Log minidisk used by SFS or CRR servers.
xxx SG1     - Storage group 1 disk used by SFS or CRR servers.
xxx SG2     - Storage group 2 disk used by SFS or CRR servers.
Y-DISK      - CMS program product disk.
3380-G32    - 3380 Cached DASD.
```

D.1.3 4381-13 (Primary System) and 4381-92E (User Side System)

```
        TYPE  VOLID   CONTENTS
        ----  -----   --------
CH 1 - 3380  PRFRES  CPRES, SPOOL, S-DISK, TDISK, OPDISK 296
       3380  PRF05   CRR CONTROL, CRR LOG1, CRR SG1, CRR SG2,CRR CRR1
       3380  MDSK01  PAGE, USERDISKS, SPOOL, TDISK, SFS CONTROL,
                     SFS SG1, SFS SG2, CRRFP1 CONTROL, CRRFP1 LOG1
                     CRRFP1 SG1
       3380  MDSK05  PAGE, USERDISKS, SPOOL, TDISK, SFS SG1, SFS SG2

CH 2 - 3380  PRF01   SPOOL, HELP, Y-DISK
       3380  PRF06   CRR LOG2, CRR CRR2,
       3380  MDSK02  PAGE, USERDISKS, SPOOL, TDISK, SFS LOG1,
                     SFS SG1, SFS SG2, CRRFP2 CONTROL, CRRFP2 LOG1
                     CRRFP2 SG1
       3380  MDSK06  PAGE, USERDISKS, SPOOL, TDISK, SFS SG1, SFS SG2

CH 3 - 3380  PRFPRO  OPDISK 295, OPDISK 395
       3380  MDSK03  PAGE, USERDISKS, SPOOL, TDISK, SFS LOG2,
                     SFS SG1, SFS SG2, CRRFP2 LOG2, CRRFP2 SG2
       3380  MDSK07  PAGE, USERDISKS, SPOOL, TDISK, SFS SG1,
                     SFS SG2, CRRFP3 CONTROL, CRRFP3 LOG1, CRRFP3 SG1
       3380  MDSK09  PAGE, USERDISKS, SPOOL, TDISK, SFS SG1, SFS SG2

CH 4 - 3380  MDSK04  PAGE, USERDISKS, SPOOL, TDISK, SFS SG1,
                     SFS SG2, CRRFP1 LOG2, CRRFP1 SG2
       3380  MDSK08  PAGE, USERDISKS, SPOOL, TDISK, SFS SG1,
                     CRRFP3 LOG2, CRRFP3 SG2
       3380  MDSK10  PAGE, USERDISKS, SPOOL, TDISK, SFS SG1, SFS SG2

CH 5 - 3088  CTC Adapter

CH 9 - 3480  TAPE
```

D.1.4 4381-13 (Resource Side System)

```
        TYPE  VOLID   CONTENTS
        ----  -----   --------
CH 1 - 3380  VMSRES  CPRES, SPOOL, S-DISK, TDISK,
       3380  VMPK01  SPOOL, Y-DISK

CH 5 - 3088  CTC Adapter

CH 9 - 3480  TAPE
```

D.1.5 9370-30

```
          TYPE  VOLID   CONTENTS
          ----  -----   --------
  CH C - 9332  VMSRES  CPRES, SPOOL, S-DISK, TDISK
         9332  MDSK01  PAGE, USERDISKS, SPOOL, TDISK, SFS CONTROL,
                       SFS SG1, SFS SG2
         9332  MDSK03  PAGE, USERDISKS, SPOOL, TDISK, SFS LOG2,
                       SFS SG1, SFS SG2

  CH D - 9332  VMPK01  OPDISKS, SPOOL, TDISK, CRR CRR2, HELP
         9332  VMPK02  Y-DISK, CRR CONTROL, CRR LOG1, CRR LOG2, CRR CRR1
                       CRR SG1, CRR SG2
         9332  MDSK02  PAGE, USERDISKS, SPOOL, TDISK, SFS LOG1,
                       SFS SG1, SFS SG2
         9332  MDSK04  PAGE, USERDISKS, SPOOL, TDISK, SFS SG1, SFS SG2
```

D.1.6 9370-80

```
          TYPE  VOLID   CONTENTS
          ----  -----   --------
  CH C - 9335  VMSRES  CPRES, SPOOL, S-DISK, TDISK
         9335  MDSK02  PAGE, USERDISKS, SPOOL, TDISK, SFS LOG1,
                       SFS SG1, SFS SG2

  CH D - 9335  VMPK01  OPDISKS, SPOOL, TDISK, CRR CRR2, HELP
         9335  MDSK03  PAGE, USERDISKS, SPOOL, TDISK, SFS LOG2,
                       SFS SG1, SFS SG2

  CH E - 9335  VMPK02  Y-DISK, CRR CONTROL, CRR LOG1, CRR LOG2, CRR CRR1
                       CRR SG1, CRR SG2
         9335  MDSK04  PAGE, USERDISKS, SPOOL, TDISK, SFS SG1, SFS SG2

  CH F - 9335  MDSK01  PAGE, USERDISKS, SPOOL, TDISK, SFS CONTROL,
                       SFS SG1, SFS SG2
```

D.1.7 3090-xxx (All DASD on the 3090 for both VM and MVS)

	TYPE	NUM	VOLID	CONTENTS
CH 1 -	3380-G32	4	PRSCnn	OFFICE
	3350	2	XXXnnn	MVS PAGING
CH 2 -	3380-G32	4	PRSCnn	OFFICE
	3350	2	XXXnnn	MVS PAGING
CH 3 -	3380-G32	4	PRSCnn	OFFICE
	3380	32	3383nn	MVS/IMS
	3350	2	XXXnnn	MVS PAGING
CH 4 -	3380-G32	4	PRSCnn	OFFICE
	3380	16	3384nn	MVS/IMS
	3350	2	XXXnnn	MVS PAGING
CH 5 -	3380	4	B1MDnn	USERDISKS
	3380	8	3385nn	MVS/IMS
	3380	8	3385nn	MVS/IPL
	3380	4	A1SWOn	SWAP HPO
	3380	8	3385nn	Maint minidisks
	3380	4	PSYS02	Staging IPL packs and minidisks
CH 6 -	3380	12	xxPGnn	PAGE
	3380	2	B1SPnn	SPOOL
	3380	2	B1TDnn	TDISK
	3380	8	3386nn	MVS/IMS
	3380	8	3386nn	MVS/IPL
CH 7 -	3380	8	x1PGnn	PAGE
	3380	2	A1SPnn	SPOOL
	3380	2	A1TDnn	TDISK
	3380	8	3387nn	MVS/IMS
	3380	8	3387nn	MVS/IPL
CH 8 -	3380	8	x1TDnn	TDISK
	3380	8	x1MDnn	USERDISKS
	3380	8	3388nn	MVS/IMS
	3380	8	3388nn	MVS/IPL
CH 9 -	3380	12	x1MDnn	USERDISKS
	3380	20	3389nn	MVS/IMS
CH A -	3380	4	x1PGnn	PAGE
	3380	4	x1SPnn	SPOOL
	3380	32	338Ann	MVS/IMS

3090-xxx (cont.)

	TYPE	NUM	VOLID	CONTENTS
CH B -	3380	4	A2SWnn	SWAP HPO
	3380	4	B2MDnn	USERDISKS
	3380	40	338Bnn	MVS/IMS
	3380	8	338Bnn	MVS/IPL
CH D -	3380	8	338nnn	MVS/IPL
CH E -	3380	8	338nnn	MVS/IPL
CH F -	3380	8	x2PGnn	PAGE
	3380	4	x2SPnn	SPOOL
	3380	4	x2TDnn	TDISK
	3380	16	338Fnn	MVS/IMS
	3380	8	338nnn	MVS/IPL
CH 11 -	3380	16	x2MDnn	USERDISKS
	3380	8	x2TDnn	TDISK
CH 12 -	3380	16	x2MDnn	USERDISKS
CH 13 -	3380	24	xxMDnn	USERDISKS
CH 14 -	3380	4	x2PGnn	PAGE
	3380	4	x2SPnn	SPOOL
CH 15 -	3380-G32	1	WKLD01	Y-DISK, S-DISK, OPDISKS, HELP
CH 16 -	3380-G32	1	WKLD01	Y-DISK, S-DISK, OPDISKS, HELP
CH 17 -	3380-G32	1	PSPT01	CPRES
	3380-G32	1	PERF2C	HPO CPRES
	3380-G32	1	WKLD02	Y-DISK, S-DISK, OPDISKS, HELP
CH 18 -	3380-G32	1	PSPT01	CPRES
	3380-G32	1	PERF2C	HPO CPRES
	3380-G32	1	WKLD02	Y-DISK, S-DISK, OPDISKS, HELP
CH 19 -	3380-G32	1	PSPT01	CPRES
	3380-G32	1	WKLD02	Y-DISK, S-DISK, OPDISKS, HELP
CH 1A -	3380-G32	1	WKLD01	Y-DISK, S-DISK, OPDISKS, HELP
CH 1C -	3380	16	AxSFnn	CRRA CONTROL, LOG1, SG1, SG2, CRR1; SFS1 CONTROL, LOG1, SG1, SG2; SFS2 CONTROL, SG1, SG2; SFS3 LOG2, SG1, SG2; SFS4 LOG1, LOG2, SG1, SG2; USERDISKS

3090-xxx (cont.)

	TYPE	NUM	VOLID	CONTENTS
CH 1D -	3380	16	AxSFnn	CRRA LOG2, CRR2; SFS1 LOG2, SG1, SG2; SFS2 LOG1, LOG2, SG1, SG2; SFS3 CONTROL, LOG1, SG1, SG2; SFS4 CONTROL, SG1, SG2; USERDISKS
CH 1E -	3380-G32	1	WKLD02	Y-DISK, S-DISK, OPDISKS, HELP
	3380	16	BxSFnn	CRRB CONTROL, LOG1, SG1, SG2, CRR1; SFS5 CONTROL, LOG1, SG1, SG2; SFS6 CONTROL, SG1, SG2; SFS7 LOG2, SG1, SG2; SFS8 LOG1, LOG2, SG1, SG2; USERDISKS
CH 1F -	3380-G32	1	WKLD01	Y-DISK, S-DISK, OPDISKS, HELP
	3380	16	BxSFnn	CRRB LOG2, CRR2; SFS5 LOG2, SG1, SG2; SFS6 LOG1, LOG2, SG1, SG2; SFS7 CONTROL, LOG1, SG1, SG2; SFS8 CONTROL, SG1, SG2; USERDISKS

D.2 Microcode Levels

The microcode level of various processors may play an important role in the overall performance. Some microcode may have various assists in them which will help improve performance while others may not. Therefore, your results may differ from our performance results just by using a different microcode level. Below is a list of each of the microcode levels (EC LEVEL) used by each of our processors.

PROCESSOR	EC LEVEL
9370-30	A28445
9370-80	364469
4381-13	A09749
4381-92E	A83520
3090	227540

D.3 Server Options

D.3.1 SFS DMSPARMS

This section gives a description of the start-up parameters used by each of the SFS servers. The start-up parameters determine the operational characteristics of the file pool server. The SFS servers use the following DMSPARMS file:

```
ADMIN MAINT OPERATOR BIGAL
FILEPOOLID SFS
NOBACKUP
NOFORMAT
USERS nnn
FULLDUMP
SAVESEGID CMSFILES
MSGS
```

For all SFS measurements, the SAVESEGID is specified to identify the segment containing the file pool server executable code. In the above example, the USERS parameter is followed by ″nnn″. This value differs for each of the processors. The SFS server configures itself with the appropriate number of user agents and buffers based on this parameter. It is recommended that USERS be set to the administrator′s best estimate of the maximum number of logged-on virtual machines that will be using the file pool during peak usage. The ratio of logged-on users to active users varies greatly on actual production machines. We assumed a 3:1 ratio and set the USERS value accordingly to 3 times the number of active users. Many of the low-end runs were done at the same time (the 70% run interval followed by the 80% run interval, etc.), so the USERS value was set to 3 times the number of active users in the final interval (90% CPU).

For more information on SFS and SFS tuning parameters, check the ***CMS Planning and Administration*** manual.

D.3.2 CRR DMSPARMS

This section gives a description of the start-up parameters 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 OPERATOR WILDBILL
NOBACKUP
FULLDUMP
SAVESEGID CMSFILES
FILEPOOLID CRRRECOV
NOACCOUNT
FORMAT
CRR
LUNAME nodeid.userid
USERS 500
CATBUFFERS 10
CTLBUFFERS 10
MSGS
REMOTE
```

D.4 Saved Segments

CMS allows the use of saved segments for shared code. The CMSFILES saved segment is required. All others are optional. Using saved segments can greatly improve performance by reducing end users′ working set sizes and thereby decreasing paging. Note that 370 and XA virtual machines can not use the same physical saved segment although logically they are the same. Before running in a different mode, the segments were resaved in the same locations. The environments in this report use the following saved segments:

- CMSFILES: Contains the SFS server code modules DMSDAC and DMSSAC, This was optional in SP 6 but is an ESA 1.0 requirement.
- HELP: Contains the CMS HELP files. Installing the HELP files in the saved segment is optional but if HELP is installed it must be in a saved segment.
- CMSINST: Holds the EXECs-in-storage. This is another optional saved segment.
- CMSVMLIB: Contains the CSL code. This is also an optional saved segment.

Measurement systems in this report also have the File Status Tables (FSTs) for the C, S, and Y-disks saved.

Appendix E. VTAM Network Definitions

E.1 VM/ESA 1.0

E.1.1 Connectivity

The following is the ATCSTRxx VTAMLST used in the AVS Performance measurements for the 512-byte runs.

ATCSTRK7

```
HOSTSA=07,
PROMPT,
NOTRACE,TYPE=VTAM,
MAXSUBA=31,
NETID=NET1,
CONFIG=K7,
SSCPID=07,
SSCPNAME=TEST,
IOBUF=(1000,512,19,,50,50),
CRPLBUF=(1200,,15,,80,80),
LFBUF=(25,,0,,10,1),
LPBUF=(400,,15,,50,50),
SFBUF=(80,,0,,50,1),
WPBUF=(2570,,0,,10,1)
```

The following is the ATCSTRxx VTAMLST used in the AVS Performance measurements for the 102400-byte runs:

ATCSTRK7

```
HOSTSA=07,
PROMPT,
NOTRACE,TYPE=VTAM,
MAXSUBA=31,
NETID=NET1,
CONFIG=K7,
SSCPID=07,
SSCPNAME=TEST,
IOBUF=(2000,4025,19,,50,50),
CRPLBUF=(200,,15,,80,80),
LFBUF=(25,,0,,10,1),
LPBUF=(100,,15,,50,50),
SFBUF=(80,,0,,50,1),
WPBUF=(200,,0,,10,1)
```

The following is the VTAMLST member used in the AVS Performance measurements to describe the 3088 link:

CTCAK7

```
CTCA07   VBUILD TYPE=CA
CTCAGRP  GROUP LNCTL=CTCA,
               DELAY=.2,
               MAXBFRU=16,
               REPLYTO=10.0
CTCALN06 LINE  ADDRESS=CD0
CTCAPU06 PU    PUTYPE=4
```

The following is the VTAMLST member used in the AVS Performance measurements to describe the AVS application:

AGWK7Q

```
AGWK7A   VBUILD TYPE=APPL
AGWT07A  APPL  AUTH=(ACQ,PASS),
               AUTHEXIT=YES,
               AUTOSES=8,
               APPC=YES,
               DMINWNL=50,
               DMINWNR=50,
               DSESLIM=100,
               MODETAB=AGWTAB,
               PARSESS=YES,
               SECACPT=CONV,
               VPACING=0
```

E.1.2 Single VTAM

The following is the DTIUSERx ASSEMBLE file used for the Single VTAM Performance measurements:

DTIUSER1

```
VMBUB    DTIGEN DTIUSER=1,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
                APPLID=VMBUB,VSAMLM=8,RPLNUM=16
         END
```

The following is the ATCSTRxx VTAMLST used for the Single VTAM Performance measurements:

ATCSTR45

```
 HOSTSA=06,
 PROMPT,
 NOTRACE,TYPE=VTAM,
 MAXSUBA=31,
 NETID=NET1,
 CONFIG=45,
 SSCPID=06,
 SSCPNAME=TEST,
 IOBUF=(1500,256,19,,50,50),
 CRPLBUF=(550,,15,,80,80),
 LFBUF=(25,,0,,10,1),
 LPBUF=(100,,15,,50,50),
 SFBUF=(80,,0,,50,1),
 WPBUF=(7710,,0,,10,1)
```

E.2 VM/ESA 1.0 370 Feature

E.2.1 Connectivity

The following is the ATCSTRxx VTAMLST used for the AVS Performance measurements on the 370 Feature:

ATCSTRY2

```
  SSCPID=4392,MAXSUBA=31,CONFIG=Y2,HOSTSA=2,
  PROMPT,DLRTCB=4,SUPP=NOSUP,NETID=NET,SSCPNAME=AGW2T,CSALIMIT=2M,
  NOTRACE,TYPE=VTAM,
  LPBUF=(64,,4,,4,22),        LARGE GENERAL PURPOSE - PAGEABLE
  LFBUF=(100,,10,,10,33),     LARGE GENERAL PURPOSE - FIXED
  WPBUF=(160,,2,F,5,5),       DEVICE CONTROL - PAGEABLE
  SFBUF=(25,,0,,1,1),         SMALL GENERAL PURPOSE - FIXED
  CRPLBUF=(80,,2,F,1,3),      RPL-COPY - PAGEABLE
  IOBUF=(800,1024,3,,10,5)    I/O BUFFERS - FIXED (NP & PP BUF REMOVED)
```

The following is the VTAMLST member used in the AVS Performance measurements to describe the 3088 link:

CTCAXX

```
CTCAAXX  VBUILD TYPE=CA
CTCA11G  GROUP  LNCTL=CTCA,ISTATUS=ACTIVE,DELAY=0,REPLYTO=25
CTCA11L  LINE   ADDRESS=A10,ISTATUS=ACTIVE,MAXBFRU=16
CTCA11P  PU     ISTATUS=ACTIVE
```

The following is the VTAMLST member used in the AVS Performance measurements to describe the AVS application:

AGW2

```
AGWAPPL2 VBUILD TYPE=APPL
AGWT02   APPL  AUTH=(ACQ,PASS),
               AUTHEXIT=YES,
               AUTOSES=8,
               APPC=YES,
               DLOGMOD=AGW2AGW3,
               DMINWNL=50,
               DMINWNR=50,
               DSESLIM=100,
               MODETAB=AGWTAB,
               PARSESS=YES,
               SECACPT=CONV,
               VPACING=0
```

E.2.2 GCS

The following is the GCS VTAMLST used for GCS measurements in this document:

GCS VTAMLST

```
SSCPID=5392,MAXSUBA=40,CONFIG=00,HOSTSA=36,
PROMPT,DLRTCB=4,SUPP=NOSUP,NETID=NETSNA,SSCPNAME=SOB,
NOTRACE,TYPE=VTAM,
LPBUF=(64,,4,,4,22),         LARGE GENERAL PURPOSE - PAGEABLE
LFBUF=(100,,10,,10,33),      LARGE GENERAL PURPOSE - FIXED
WPBUF=(160,,2,F,5,5),        DEVICE CONTROL - PAGEABLE
SFBUF=(25,,0,,1,1),          SMALL GENERAL PURPOSE - FIXED
CRPLBUF=(80,,2,F,1,3),       RPL-COPY - PAGEABLE
IOBUF=(200,1024,3,,10,5)     I/O BUFFERS - FIXED (NP & PP BUF REMOVED)
```

Glossary of Performance Terms

This glossary uses postscripts to reflect the sources of the data described in this document. 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:

(CONN): Refers to connectivity data.

(H): Highlights data that originates in an internal use only hardware monitor.

(Q): Denotes data from the SFS QUERY FILEPOOL STATUS command. This command requires administrator authority.

(R): Marks data from the IBM Internal Use Only tool RESPONS2.

(RTM): Specifies data originating from the program product VM / Extended Architecture Realtime Monitor / Systems Facility (RTM).

(Server): Indicates that the data is for specific virtual machines. (e.g. SFS, CRR, or VTAM) If there is more than one virtual machine of the same type, these data fields are for all the virtual machines of that type.

(T): Identifies data from the program product Teleprocessing Network Simulator (TPNS).

(VMMAP): Indicates the data originates in the program product Virtual Machine Facility / 370 Performance / Monitor Analysis (VMMAP).

(VMPRF): Denotes data from the program product VM Performance Reporting Facility (VMPRF). VMPRF reduces XA MONITOR data.

ACTIVE USERS: The number of users executing commands during a measurement.

Agent: The unit of subdispatching within a CRR or SFS filepool server.

Agents Held: The average number of agents that are in a Logical Unit of Work (LUW). This is calculated by:

((The sum of Agent Holding Time over all production filepools) / 1000) / (the sum of SFSTIME over all production filepools).

Agent Holding Time is from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

Agents In Call: The average number of agents that are currently processing SFS server requests. This is calculated by:

((The sum of Filepool Request Service Time over all production filepools) / 1000) / (The sum of SFSTIME over all production filepools).

Filepool Request Service Time is from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

AVE TOT (R): The average total response time for all CMS commands executed. The total response time is calculated from the console read of the input line to the next issuance of a console read. This is computed by the user written program RESPONS2.

AVE TOT (T): The average total response time for all commands executed from the time the message is sent until the last recieve prior to the next send. This includes MOREs from the host to the simulated user (TPNS) and CLEARS from TPNS to the host.

AVG FIRST (T): See First Response (T)

AVG LAST (T): See Last Response (T)

Bactrian: A two-humped curve used to represent the think times for both active users and users who are logged on but inactive.

Blocking Factor (Blocks/BIO): The average number of blocks read or written per Block I/O Request. This is calculated by:

(The sum of Total DASD Block Transfers over all production filepools) / (The sum of Total BIO Requests over all production filepools).

Total DASD Block Transfers and Total BIO Requests are from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

Chaining Factor (Blocks/IO): The average number of blocks read or written per I/O Request. This is calculated by:

(The sum of Total DASD Block Transfers over all production filepools) / (The sum of Total I/O Requests over all production filepools).

Total DASD Block Transfers and Total I/O Requests are from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

Checkpoint: 1) In an SFS filepool server, the periodic processing that records a consistent state of the filepool 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. This allows areas of the logs to be reused.

Checkpoint Duration: The time, in seconds, required to process an SFS checkpoint. This is calculated by:

((The sum of Checkpoint Time over all production filepools) / 1000)) / (The sum of Checkpoints Taken over all production filepools).

Checkpoint Time and Checkpoints Taken are from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

Checkpoint Rate: Count of CRR checkpoints performed per second.

Checkpoint Utilization: The percentage of time an SFS filepool server spends performing checkpoints. This is calculated by:

((The sum of Checkpoint Time over all production filepools) / 1000) * 100 / (The sum of SFSTIME over all production filepools).

Checkpoint Time is from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

CMDS/SEC (CONN): In the connectivity environment, the number of APPC SEND/RECEIVE pairs completed per second for the end users.

CMDS/SEC (R): The command rate as determined by RESPONS2 where a command is any user interaction with the system.

Command (R): See Transaction (R).

Command (RTM): See Transaction (RTM).

Command (T): See Transaction (T).

Coordinated Resource Recovery (CRR): VM implementation of the SNA LU 6.2 Syncpoint architecture. Allows multiple resources to be updated and work to be committed together.

Coupled Commit: The process used by SFS to combine a filepool request with a commit request. The Syncpoint Manager is shoulder tapped before and after the combined requests are made.

CPCPU (VMMAP): The total CP CPU time as a percentage of elapsed time.

CP/CMD (CONN): The average number of CP CPU seconds consumed per APPC SEND/RECEIVE pair for the user virtual machine. This is calculated by:

(CPCPU (VMMAP) / 100) / (The total number of APPC SEND/RECEIVE pairs).

CP/CMD (H): The average amount of CP CPU time used per command in milliseconds. This is calculated by:

(SUP / ETR (T)).

SUP is from an internal use only hardware monitor and ETR (T) is from TPNS.

CP/CMD (R): The average user CP CPU per command, excluding uncompleted commands, service machines, and system overhead not charged to a specific virtual machine. This is the difference between PBT/CMD (R) and VIRT/CMD (R).

CP/CMD (Server): CP CPU time executed in the designated server machine per command. This is:

(TOTCPU (Server) - VIRTCPU (Server)) / TOTAL CMDS

TOTAL CMDS is from RESPONS2 or TPNS. If there is more than one server virtual machine of the same type this is the sum of CP CPU times for all like server machines.

CP/CMD (VMMAP): The average number of CP CPU seconds consumed per command. This is calculated by:

(CPCPU (VMMAP) / 100) / CMDS/SEC (R).

CP/ESA: The Control Program (CP) component of VM/ESA for large systems. This is based on the VM/XA 2.1 CP component.

CP/370: The Control Program (CP) component of the VM/ESA 370 Feature. This is based on the VM/SP 6 CP component.

Deadlocks: The total number of SFS filepool deadlocks that occurred during the measurement interval summed over all production filepools. A deadlock occurs when two users each request a resource that the other currently owns.

The values for Deadlocks are obtained from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

DIAG/CMD (RTM): The number of DIAGNOSE instructions executed per command. This is:

(The sum of all DIAGNOSE instructions on the RTM PRIVOPS screen) / (The sum of all transactions over all TPNSs).

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

(The DIAG 08 rate on the RTM PRIVOPS screen) / (ETR (T)).

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

(The DIAG 10 rate on the RTM PRIVOPS screen) / (ETR(T)).

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

(The DIAG 58 rate on the RTM PRIVOPS screen) / (ETR(T)).

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

(The DIAG A4 rate on the RTM PRIVOPS screen) / (ETR(T)).

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

(The DIAG A8 rate on the RTM PRIVOPS screen) / (ETR(T)).

DIAGNOSE X′98′: An option of the CP DIAGNOSE command. This allows a specified virtual machine to lock and unlock virtual pages and to execute its own channel program. See the ***VM/ESA System Facilities for Programming*** for more information.

DMSDAC: One of the two modules containing SFS filepool server code.

DMSSAC: One of the two modules containing SFS filepool server code.

DYNAMIC (VMMAP): Storage allocated to virtual machine pages and pageable nucleus pages. This value is obtained by dividing the VMMAP parameter DYNAMIC PAGING AREA by 1024.

EDF: Enhanced Disk Format. This refers to the CMS minidisk file system.

EMUL ITR (RTM): 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 (H): The amount of processor time spent in emulation mode per command in milliseconds. This is:

(PROB / ETR (T)).

PROB is from an internal use only hardware monitor and ETR (T) is from TPNS.

ESA: VM/Enterprise Systems Architecture. The latest release of VM combining features of VM/HPO 5, VM/SP 6, VM/XA 2.0, and VM/XA 2.1.

ETR: External Throughput Rate. The number of units of work completed per unit of elapsed time.

ETR (R): The number of CMS commands executed per second, computed by RESPONS2.

ETR (T): The total number of commands executed per second, computed by TPNS. This is the sum over TPNS reports of TOTAL XMITS by TPNS divided by the elapsed time calculated from TPNS reports.

Expanded Storage: An optional integrated high-speed storage facility, available on certain processors, that allows for the rapid transfer of 4K blocks between itself and real storage.

Extended Storage: Processor storage ranging from 16M to 64M.

FAVOR: The CP SET FAVOR command. It indicates to CP that the virtual machine is always runnable. This means that it goes directly to the dispatcher when it has work to do. See the ***CP System Command Reference*** for more details.

FAVOR 100: The CP SET FAVOR command. When FAVOR is set at 100 the scheduler adjusts the virtual machine′s dispatching priority so it tends to be placed at the top of the dispatch list. Note that specifying 100% does not result in 100% of processor time being as-

signed to the virtual machine. See the ***CP System Command Reference*** for more details.

Filepool: In SFS, a collection of minidisks managed by a server machine.

FILE IO/CMD (Q): The total number of CMS file system I/Os executed per commnad. This includes both EDF and SFS I/Os. This is calculated by

(MDSK/CMD) / (IO/CMD).

MDSK/CMD is from RESPONS2 while IO/CMD is from the QUERY FILEPOOL STATUS command.

FP REQ/CMD (Q): Total file pool requests per command. This is calculated by:

Total File Pool Requests / TOTAL CMDS

TOTAL CMDS is from RESPONS2 or TPNS. Total File Pool Requests is from the QUERY FILEPOOL STATUS command.

FIRST RESPONSE (R): The average response time of all CMS commands executed. The response time is calculated from console read of the input line to the console write of an output line or CMS ready message. This is computed by the user written program RESPONS2. This represents the internal response time and doesn't include line time.

FIRST RESPONSE (T): The first reply that returns to the screen. For non-fullscreen commands this is the command reflect on the screen. These include line delay.

FREE (VMMAP): Fixed storage set aside for CP control blocks. This value is obtained by dividing the VMMAP parameter FREE STORAGE SIZE by 1024.

FSID: Full Screen Internal Driver. An IBM Internal Use Only tool used to simulate locally attached full screen terminals (3277's).

G: Gigabytes. 1024 megabytes.

HPO: VM/High Performance Option.

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

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

The sum of Total I/O Requests over all production filepools / TOTAL CMDS (Q).

Total I/O Requests is from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

IO TIME/CMD (Q): Total elasped time spent doing SFS file I/O's per command. This is calculated by:

The sum of Total BIO Request Time over all production filepools / TOTAL CMDS (Q).

Total BIO Request Time is from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

ITR: Internal Throughput Rate. This is the number of units of work accomplished per unit of CPU busy time in an unconstrained environment.

ITR (H): See ITR. This is calculated by:

(ETR (T)) / (Total CPU Util/100)

where Total CPU Util is from an internal use only hardware monitor.

ITR (RTM): This is the average number of transactions per second of CPU time. This is found in the SYSITR field on the TRANSACT screen.

ITR (VMMAP): See ITR. For VMMAP data in this report this is the inverse of PBT/CMD. For measurements involving paging constraints this number doesn't reflect the true definition of an ITR but is still a useful measure of throughput.

ITRR (H): Internal Throughput Rate Ratio. This is the ITR (H) normalized to a specific run.

ITRR (RTM): Internal Throughput Rate Ratio. This is the ITR (RTM) normalized to a specific run.

K: Kilobytes. One kilobyte is 1024 bytes.

LAST RESPONSE (T): The average response time for the last response to the screen.

Limp Mode: In a CRR environment, when the CRR recovery server is unavailable.

LOG I/O RATE: For CRR, the rate per second of physical I/O's for logging.

LOG I/O TIME: For CRR, average elapsed time for log I/O processing in milliseconds.

Log Ring: The portion of CRR log disks where actual syncpoint data is recorded. This is a wrappable area that is maintained by the checkpoint process.

LOG REQ RATE: The number of log write requests made of the CRR Recovery Server per second.

LPA: Link Pack Area.

LSQA: Logical System Queue Area.

LUW Rollbacks: The total number of SFS logical units of work that were backed out during the measurement interval summed over all production filepools.

The values for LUW Rollbacks are obtained from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

M: Megabytes. One megabyte is 1,048,756 bytes.

MDC Hit Ratio: Minidisk Cache Hit Ratio. The ratio of requested blocks that are found in extended storage for a read operation to the total number of read requests. This is the from the RTM MDHR field on the SYSTEM screen.

MDC Mods: Minidisk Cache Mods. The number of times CMS Blocks were written in the cache, excluding the writes that occured as a result of minidisk cache misses. This is from the RTM MDC_MO field on the SYSTEM screen.

MDC Reads: Minidisk Cache Reads. The number of times CMS blocks were found in the cache as the result of a read operation. This is the from the RTM MDC_HT field on the SYSTEM screen.

MDC Writes: Minidisk Cache Writes. The number of CMS Blocks moved from main storage to extended storage. This is the from the RTM MDC_PW field on the SYSTEM screen.

MDSK/CMD (R): The average number of virtual DASD I/Os issued by the EDF file system per command. This number excludes uncompleted commands and service machine activities.

Millisecond: One one-thousandth of a second.

MPG: Multiple Preferred Guests. Allows multiple guests to be run as V=F. This requires the PR/SM microcode.

Msec per BIO Request: Average time required to process a block I/O request in milliseconds. This is calculated by:

(The sum of Total BIO Request Time over all production filepools) / (The sum of Total BIO Requests over all production filepools).

Total BIO Request Time and Total BIO Requests are from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

Multi-user: Refers to a connectivity environment with many users making requests of one or more resources.

NON SP / CMD (VMMAP): The average number of nonspooled I/Os per command for the system or the given virtual machine.

NONTRIV INT: Non-trivial Internal response time. The number of transactions that completed with more than one drop from Q1 and/or one or more drops from Q0, Q2, or Q3 per second. This is from TOTALTTM for the RTM NTRIV field on the TRANSACT screen.

NUCLEUS (VMMAP): The part of CP that is resident in main storage. This value is obtained by dividing the VMMAP parameter NUCLEUS SIZE by 1024.

One-phase Commit: Syncpoint processing which results in SPM driving all resource adapters just once. This is also known as a simple commit.

PAGE RATE (VMMAP): The number of pages per second moving into and out of main storage.

PAGE/CMD (CONN): The average number of pages moving in and out of storage per APPC SEND/RECEIVE pair for the system or the given virtual machine. This is calculated by:

(PAGE RATE (VMMAP)) / (The total number of APPC SEND/RECEIVE pairs).

PAGE/CMD (R): Average user page reads for completed commands. This number excludes uncompleted commands and service machine activities.

PAGE/CMD (RTM): The number of DASD page I/Os per command. This is calculated by:

(READ/SEC + WRITE/SEC) / ETR (T)

where READ/SEC and WRITE/SEC are from RTM.

PAGE/CMD (Server): Page I/Os in the server per command. It is calculated by the following formula:

(PAGE READS + PAGE WRITES)*1000 / TOTAL CMDS.

PAGE READS and PAGE WRITES are from VMMAP. TOTAL CMDS is from RESPONS2. If there is more than one server of the same type this is summed across all servers of that type.

PAGE/CMD (VMMAP): The average number of pages moving in and out of storage per command for the system or the given virtual machine. This is calculated by:

(PAGE RATE (VMMAP)) / CMDS/SEC (R)

PBT/CMD (CONN): Processor Busy Time per APPC SEND/RECEIVE pair. This is calculated by:

(TOTCPU / 100) / (The total number of APPC SEND/RECEIVE pairs).

TOTCPU is from VMMAP.

PBT/CMD (H): This is the number of milliseconds of CPU activity per the average number of transactions completed. This is calculated by:

(CPUT) / ETR (T)

CPUT comes from an internal use only hardware monitor. ETR (T) is from TPNS.

PBT/CMD (R): Average user CPU time per command. This number excludes uncompleted commands, service machine activities, and system time.

PBT/CMD (Server): Total CPU time executed in the server machine per command. This is calculated by:

TOTCPU (Server) / TOTAL CMDS.

TOTAL CMDS is from TPNS for CP/ESA measurements and RESPONS2 for CP/370 runs.

PBT/CMD (VMMAP): Processor Busy Time per command in seconds. This is calculated by:

(TOTCPU / 100) / CMDS/SEC (R).

TOTCPU is from VMMAP and CMDS/SEC is from RESPONS2.

PCA: See Protected Conversation Adapter.

PCTCPUQ (VMMAP): The average percent of active users waiting for CPU.

PCTIOQ (VMMAP): The average percent of active users waiting for I/O.

PCTPAGQ (VMMAP): The average percent of active users waiting for paging.

PCTSTGQ (VMMAP): The average percent of active users waiting for storage.

PER: The CP PER command (Program Event Recording). The PER command allows tracing of instructions, branches, storage altering instructions, and general registers.

PGBLPGS (RTM): The number of system pageable pages available. This is from the PPAG field on the RTM SYSTEM screen.

PGBLPGS (VMMAP): The number of system pageable pages available.

Priviledged Operation: Any 370 instruction that must be executed in supervisor state.

PRIVOP RATE (VMMAP): The number of virtual machine privileged instructions simulated by CP per second. This does not include those instructions processed by the VMA (Virtual Machine Assist) feature. This is calculated by:

(PRIVOP RATE (VMMAP)) / (CMDS/SEC (R))

Note: PRIVOPS are recorded differently in 370 and XA modes.

PRIVOP/CMD (CONN): The number of virtual machine privileged instructions simulated per APPC SEND/RECEIVE pair. This does not include those instructions processed by the VMA (Virtual Machine Assist) feature. This is:

PRIVOP RATE / (The total number of APPC SEND/RECEIVE pairs).

PRIVOP RATE is from VMMAP. **Note**: PRIVOPS are recorded differently in 370 and XA modes.

PRIVOP/CMD (RTM): The number of virtual machine privileged instructions simulated per command. This does not include DIAGNOSE instructions. This is:

(The sum of the priviledged operations listed on the RTM PRIVOPS screen) / (The sum of all transactions over al TPNSs).

Note: PRIVOPS are recorded differently in 370 and XA modes.

PRIVOP/CMD (VMMAP): The number of virtual machine privileged instructions simulated per command. This does not include those instructions processed by the VMA (Virtual Machine Assist) feature. This does include DIAGNOSE instructions. This is:

PRIVOP RATE / CMDS/SEC (R).

PRIVOP RATE is from VMMAP and CMDS/SEC is from RESPONS2. **Note**: PRIVOPS are recorded differently in 370 and XA modes.

PROCESSORS: The data field denoting the number of processors within a 3090-600J that were active during a measurement. This is from the NC field under CPU statistics on the RTM GENERAL screen.

Production Filepool: An SFS filepool in which users are enrolled with space. All SFS read/write activity is to production filepools.

Protected Conversation Adapter (PCA): CMS code which runs in each end user virtual machine. The PCA is the interface between the SPM and the protected conversation.

Protected Conversation: An APPC conversation established with the SYNCLEVEL = SYNCPT parameter. These conversations can participate in CRR processing.

PR/SM: Processor Resource/Systems Manager Feature. Hardware feature that supports VM V=F guests and the logical partitioning of a 3090 processor complex.

QDROP OFF: A version of the CP SET QDROP command. This keeps the user's page and segment tables from being scanned. It also eliminates the placing of the user's resident pages on the flush list.

QDROP OFF USERS: A version of the CP SET QDROP command. This gives a temporary extension to users connected with a QDROP OFF service machine causing them to be treated as QDROP OFF virtual machines while they are communicating with the service machine. This applies to all virtual machines communicating with VMCF, IUCV or APPC/VM.

QUICKDSP ON: QUICKDSP can be specified either via a CP command or in the CP directory entry. 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.

READS/SEC (RTM): The number of read operations done for system paging per second. For RTM, this is the PAGREAD field of the SYSTEM screen divided by elapsed time.

REAL (RTM): The actual real storage size of the CPU in megabytes. This is nnnM on the RTM GENERAL screen header.

Recovery Server: The CRR server machine. There is one and only one per system and is responsible for logging coordination information.

Regression: Measurements that compare the current and previous releases where the only changed variable is the release level.

Relative Share: For CP/ESA, determines the virtual machine's scheduling priority.

RELSHARE: The virtual machine's relative share.

Resource: A virtual machine that handles user requests. A resource can handle requests from multiple users.

Resource Adapter: The interface between a CRR resouce manager and the Syncpoint Manager (SPM).

Resource Manager: Component that controls a protected resource. The resource manager communicates to the SPM through a resource adapter. For SFS, the filepool server is the resource manager.

RESPONS2 (R): IBM Internal Use Only Monitor reduction program. Reduces RESPONSE and SCHEDULE class monitor records created by the VM/370 MONITOR and produces reports on response times and resource utilization for the commands found in the monitor data.

Rollbacks Due to Deadlock: The total number of LUW rollbacks due to deadlock that occurred during the measurement interval over all production filepools. A rollback occurs whenever a deadlock condition cannot be resolved by the SFS server.

The values for Rollbacks Due to Deadlock are obtained from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

RTM: VM / Extended Architecture Realtime Monitor / Systems Facility. A Program Product realtime monitor and diagnostic tool for performance monitoring, analysis, and problem solving.

RTM/SF: See RTM.

SDLC: Synchronous Data Link Control. A discipline for managing syncronous, code-tranparent, serial-by-bit information transfer over a link.

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

(The sum of SFSTIME over all production filepools) / (The sum of Checkpoints Taken over all production filepools).

Checkpoints Taken is from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

SFSTIME: The elapsed time between QUERY FILEPOOL STATUS invocations done at the beginning and end of the measurement.

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

(File Pool Request Service Time) / (TOTAL CMDS (Q)).

File Pool Request Service Time is from the QUERY FILEPOOL STATUS command. **Note**: For 370 Feature measurements there is only one filepool.

SHRPGS (RTM): The number of shared frames currently resident. This is from the SHAR field on the RTM SYSTEM screen.

SHRPGS (VMMAP): The number of pages being shared among users.

SIE: XA Architecture instruction to Start Interpretive Execution. Instruction used to run a virtual machine in emulation mode.

Simple Commit: See One-phase Commit.

Single-thread: Refers to a connectivity environment with one user making requests of one resource.

SPM: See Syncpoint Manager.

Special Operations: Instructions that cause either CP or the hardware microcode to provide services for the virtual machine. This includes priviledged instructions, normally invalid opcodes (i.e. DIAG), and the first time a non-shared page is referenced.

SRM: System Resource Manager.

SQA: System Queue Area.

STARS: System Trace Analysis Reports. An IBM Internal Use Only program package used to provide information on instruction path length and storage references while executing given functions in a VM environment.

Syncpoint: The process by which work with resources participating in CRR is either all committed or rolled back (backed out).

Syncpoint Manager (SPM): The portion of CRR that runs in the end user's virtual machine.

SYNCPT RATE: The number of times work was committed per second through the two-phase commit process.

System Time: Processor time executed in CP that CP can not apportion to any given virtual machine.

TCB: Task Control Block.

TOT CPU/CMD (CONN): The average number of CPU seconds per APPC SEND/RECEIVE pair required by the system or the given virtual machine.

TOT CPU/CMD (VMMAP): The average number of CPU seconds per command required by the system or the given virtual machine.

TOT CPU/CMD (VMPRF): The average number of CPU seconds per command required by the system or the given virtual machine.

TOT INT: Total Internal Response Time. The rate per second of all trivial and non-trival transactions. This is the value for TOTALTTM for ALL_TRANS on the RTM TRANSACT screen.

TOTAL (H): The total CPU Utilization for a given measurement for all processors. This is calculated by

(TOT CPU UT) * (Number of CPUs).

TOT CPU UT is from an internal use only hardware monitor.

TOTAL CMDS (Q): The total number of commands issued during the QUERY FILEPOOL STATUS (QFP Status) measurement interval. This is calculated by:

(Length of the QFP Status measurement interval) / (CMDS/SEC).

CMDS/SEC is from TPNS in the CP/ESA environment and RESPONS2 for CP/370.

TOTAL CMDS (R): The total number of RESPONS2 commands completed in a measurement interval.

TOTAL CMDS (VMMAP): The total number of commands issued during the VMMAP measurement interval. This is calculated by:

(Length of the VMMAP measurement interval) * (CMDS/SEC (R)).

TOTAL STORAGE (VMMAP): The amount of real storage available. This value is obtained by the VMMAP parameter STORAGE TOTAL by 1024.

TOTCPU (Server): The total amount of CPU time in the server virtual machine. This is from VMPRF for CP/ESA environments and from VMMAP for CP/370. If there is more than one of the same type of server machine, this is summed over all of them.

TOTCPU (VMMAP): The total processor busy time as a percentage of elapsed time. This

includes all time charged to the virtual machine but does not include system time.

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

(%TOTCPU * %Emulation * Number of processors) / 100

All values are from an internal hardware monitor.

TPNS: Teleprocessing Network Simulator. A program product terminal and network simulation tool which provides system performance and response time information.

TRACE: 1) An area of fixed real storage in which critical operations of CP are recorded for error detection and correction purposes. This value is obtained by dividing the VMMAP parameter INTERNAL TRACE TABLE by 1024. 2) A recording of all activity within a virtual machine.

Transaction (R): A RESPONS2 transaction. RESPONS2 considers every terminal input to be a command, even though it may be a subcommand or a reply to a query. This does not include those commands that do not produce RESPONSE class MONITOR records. These are commands issued within EXECs and some CP commands.

Transaction (RTM): For a single-user virtual machine a transaction should roughly correspond to a CMS 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.

Transaction (T): This is the interval from the time the command is issued until the last recieve prior to the next send. This includes clear screens as a result of an intervening more holding condition.

TRIV INT (RTM): Trivial Internal Response Time. The rate per second for transactions that complete with one and only one drop from Q1 and no drops from Q0, Q2, and Q3. This is from TOTALTTM for the TRIV field on the RTM TRANSACT screen.

TRIV RESP (VMMAP): The average response time of commands that finish within one Q1 time slice. This is taken from the VMMAP report.

TRREAD: Trace Read. An IBM Internal Use Only program in the STARS package used to read VM trace data from the spool into CMS files.

Two-phase Commit: Syncpoint process which results in the SPM driving resource adapters twice, first to determine if all resources can commit and second in order to commit or roll back based on replies from first phase.

TVR (H): Total to Virtual Ratio. This is the ratio of total CPU utilization to virtual CPU utilization as reported by an internal use only hardware monitor This is:

%CPU / %EM.

%CPU and %EM are fields on the RTM GENERAL screen.

UNIQUE SHARED: Number of unique shared pages touched across the commands traced.

User: 1) A virtual machine associated with a simulated end user. 2) In the connectivity environment, a virtual machine that makes requests of a resource virtual machine. Each user can connect to only one resource.

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 (H): CPU Utilization per processor. This is the TOT CPU UT field from an internal use only hardware monitor.

VIO RATE (RTM): The total number of all virtual I/O requests per second for all users in the system. This is from the ISEC field on the RTM SYSTEM screen.

VIO RATE (VMMAP): The number of virtual I/O requests simulated per second.

VIO/CMD (RTM): The average number of virtual I/O requests per command for all users in the system. This is calculated by:

ISEC / TOTALCNT of ALL_TRANS.

ISEC and TOTALCNT of ALL_TRANS are fields on the RTM SYSTEM and TRANSACT screens.

VIO/CMD (CONN): Average number of simulated virtual I/O requests per APPC SEND/RECEIVE pair.

VIO/CMD (VMMAP): Average number of simulated virtual I/O requests per command.

VIRTCPU (Server): Virtual processor time in the server virtual machine. This is from VMPRF for CP/ESA environments and from RESPONS2 for CP/370. If there is more than one of the same type of server machine, this is summed over all of them.

VIRTCPU (VMMAP): The total problem state CPU time as a percentage of elapsed time.

VIRT/CMD (R): Average user virtual CPU time per command in seconds. This number excludes uncompleted commands and service machine activities.

VIRT/CMD (Server): Virtual CPU time executed in the server machine per command. This is:

VIRTCPU (Server) / TOTAL CMDS.

TOTAL CMDS is from TPNS in the CP/ESA environment and RESPONS2 for CP/370.

VIRT/CMD (VMMAP): The average number of problem state CPU seconds consumed per command. This is calculated by:

(VIRTCPU / 100) / CMDS/SEC.

VIRTCPU is from VMMAP and CMDS/SEC is from RESPONS2.

VMMAP: Virtual Machine Facility / 370 Performance / Monitor Analysis. A Program Product which uses data collected by the VM/370 MONITOR and produces reports summarizing the performance and utilization of the running system.

VMPRF: VM Performance Reporting Facility. A Program Product which produces performance reports and history files from VM/XA MONITOR data.

VPGM INT RATE: The number of virtual program checks simulated per second.

VPGM INT/CMD: The number of virtual program checks simulated per command.

VSCSs: The number of virtual machines running VSCS during a measurement interval.

VTAMs: The number of virtual machines running VTAM during a measurement interval.

V=F: Virtual equals fixed machine. A virtual machine which 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 to this area. Requires PR/SM hardware feature to be installed.

V=R: Virtual equals real machine. Virtual machine which has fixed, contiguous area of real storage starting at page 0. CP does not page to 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 (Server): Average working set of the server virtual machine. Taken directly from the VMMAP user resource summary. If there is more than one of the same type of server machine, this is summed over all of them.

WKSET (VMMAP): The average working set size over all active virtual machines in the system, expressed in K bytes. This is taken from VMMAP's Monitor Statistical Summary Report. When provided for a given virtual machine, this is the average working set for that virtual machine and comes from VMMAP's User Resource Utilization Summary Part 1.

WKSET (VMPRF): 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 UCLASS_RESOURCE_UTIL report.

WRITES/SEC (RTM): The number of write operations done for system paging. This is from the PAGWRIT field of the RTM SYSTEM screen.

XA: VM/Extended Architecture.

XSTOR: Expanded Storage. An optional hardware high-speed storage facility. Fast transfer of 4K pages is possible between real storage and expanded storage. CP uses expanded storage for system paging, data caching, and mindisk caching.

XSTOR (RTM): The amount of expanded storage available on a processor for a specified measurement.

XSTOR IN/SEC (RTM): The number of pages read into main storage from expanded storage. This includes fastpath and non-fastpath pages. The value is the sum of NSECs for the XST_PGIS and XST_PGIF fields from the RTM SYSTEM screen.

XSTOR OUT/SEC (RTM): The number of pages written from main storage into expanded storage. This is NSEC for the XST_PGO field of the RTM SYSTEM screen.

XSTOR/CMD (RTM): The number of pages read into main storage from expanded storage

and written to expanded storage from main storage per command. This is calculated by:

(XSTOR IN/SEC + XSTOR OUT/SEC) / ETR (T).

XSTOR IN/SEC and XSTOR OUT/SEC are from RTM.

1ST RESP (R): See FIRST REPONSE (R).

1ST RESP (T): See FIRST REPONSE (T).

Related Publications

The following is a list of related publications:

- ***Connectivity Driver Reference,*** Jennifer L. Thrall, VM Performance, G21/17-2, Endicott, NY, September, 1987
- ***CRRTOOL User's Guide and Reference,*** William J. Bitner, VM Performance, G25/17-2, Endicott, NY, June 5, 1989. VM/XA System Performance Evaluation, 46E/990, Kingston, NY,
- ***Full Screen Internal Driver Version 2,*** William J. Bitner, VM Performance, G25/17-2, Endicott, NY, May 6, 1988.
- ***IBM System/370 Principles of Operation*** GA22-7000.
- ***IBM ES/3900 Model S VM/XA SP Release 2 Capacity Study,*** ZZ05-0459.
- ***RESPONS2 Full Screen Response Time Reduction Program User Guide and Reference,*** R. J. Schultz - STD, December 18, 1986, TR 01.A098.
- ***Results of Running Two VM/VTAM's on an HPO 5.0 System,*** Dave Lewis, VM/XA System Performance Evaluation, 46E/990, Kingston, NY, May 4, 1987.
- ***Revised CMS Intensive Workload (FS7B),*** Allan B. Lebovitz, VM Performance, G25/17-2, Endicott, NY, January 4, 1990.
- ***STARS User's Guide,*** Gary A. Hine, VM Performance, G21/17-2, Endicott, NY, November 25, 1986.
- ***Tuning Your VTAM Service Machine,*** GG66-0235.
- ***TPNS User's Guide,*** SH20-2490.
- ***Virtual Machine / Extended Architecture / Realtime Monitor / Systems Facility (RTM),*** SH26-7000.
- ***Virtual Machine Monitor Analysis Program (VMMAP),*** SC34-2166.
- ***VM Performance Reporting Facility (VMPRF),*** SC23-0460.
- ***VM/ESA 1.0 CMS Administration Reference,*** SC24-5446.
- ***VM/ESA 1.0 CMS Planning and Administration,*** SC24-5445.
- ***VM/ESA 1.0 CP Planning and Administration for 370,*** SC24-5430.
- ***VM/ESA 1.0 CP Programming Services,*** SC24-5520.
- ***VM/ESA 1.0 CP Programming Services for 370,*** SC24-5435.
- ***VM/ESA 1.0 CP System Command Reference for 370,*** SC24-5434.
- ***VM/SP HPO 5.0 Performance Study Report,*** Phil Coons et al., VM/XA System Performance Evaluation, 46E/990, Kingston, NY, June, 1987.
- ***VM/SP 6 Performance Report,*** ZZ05-0457.
- ***VM/SP 6 Shared File System Performance Report,*** ZZ05-0458.
- ***VM/XA SP 2.0 Performance Reference,*** ZZ20-9999.
- ***VM/XA SP 2.1 Performance Report,*** Brian Kettler, VM/XA System Performance Evaluation, 46E/990, Kingston, NY,

Heading ID's

id	File	Page	Heading References
i	VM10RPTF	1	Introduction
g	VM10RPTF	2	General Observations
ge	VM10RPTF	2	VM/ESA 1.0
genobs	VM10RPTF	2	Comparisons to Other Workloads
ger	VM10RPTF	3	Regression
gex	VM10RPTF	10	Comparison of 370 and XA Modes
ges	VM10RPTF	13	EDF and SFS Comparisons 29
gec	VM10RPTF	14	Connectivity 29
gecap	VM10RPTF	14	APPC/VM
gecav	VM10RPTF	18	AVS
gep	VM10RPTF	18	Processor Capacity
gemi	VM10RPTF	21	Migration 25
gemv	VM10RPTF	25	Single VTAM
gemmvs	VM10RPTF	26	MVS Guest
g3	VM10RPTF	28	VM/ESA 1.0 370 Feature 32
g3r	VM10RPTF	28	CMS Regression
g3s	VM10RPTF	28	EDF and SFS Comparisons
g3c	VM10RPTF	29	Connectivity
g3cap	VM10RPTF	30	APPC/VM
g3cav	VM10RPTF	31	AVS
g3ct	VM10RPTF	32	TSAF
g3cr	VM10RPTF	32	Coordinated Resource Recovery 131
g3cr1	VM10RPTF	33	Regression Performance
g3cr2	VM10RPTF	33	New Function Performance
g3g	VM10RPTF	34	GCS
g316	VM10RPTF	34	Greater Than 16M Real Storage
se	VM10RPTF	36	Specific Measurements: VM/ESA 1.0
sei	VM10RPTF	36	Introduction
ser	VM10RPTF	38	Regression
serch	VM10RPTF	38	CMS Intensive Measurements 4
serc20	VM10RPTF	41	3090-600J / 512M / VM/XA SP 2.0 8
serc211	VM10RPTF	43	3090-200J / 256M / VM/XA SP 2.1 6
serc212	VM10RPTF	45	3090-600J / 512M / VM/XA SP 2.1 8
semv	VM10RPTF	47	Comparison of 370 and XA Modes
se3x	VM10RPTF	47	370 to XA Measurements 10, 21
ses	VM10RPTF	50	EDF and SFS Comparisons 13
ses20	VM10RPTF	50	Introduction
ses30	VM10RPTF	51	Measurements
ses40	VM10RPTF	51	3090-600J / 512M / EDF and SFS at Equal Utilization
ses50	VM10RPTF	55	3090-600J / 512M / EDF and SFS with Equal Users 52
sec	VM10RPTF	59	Connectivity
seci	VM10RPTF	59	Introduction
secm	VM10RPTF	59	Local APPC Measurements
secap	VM10RPTF	59	APPC/VM Single-user 14, 15, 16
secapm	VM10RPTF	62	APPC/VM Multi-user

Figure ID′s

FGEPCR	VM10RPTF	20	13:	
FGEMITC	VM10RPTF	22	14:	
FGEMITR	VM10RPTF	23	15:	
FGEMITI	VM10RPTF	24	16:	
FGEMV	VM10RPTF	27	17:	
fg3cap1	VM10RPTF	30	18:	30
fg3cav1	VM10RPTF	31	19:	31
fg3ct1	VM10RPTF	32	20:	32
fg3cr1	VM10RPTF	34	21:	34

Table ID's

id	**File**	**Page**	**Table References**	
secapt1	VM10RPTF	61	8:	60
secapt2	VM10RPTF	63	9:	62
secavt1	VM10RPTF	65	10:	64
singv	VM10RPTF	74	14:	73
multv	VM10RPTF	77	15:	73
s3cpt1	VM10RPTF	108	27:	
s3cpt2	VM10RPTF	110	28:	
s3cpt3	VM10RPTF	112	29:	
s3cavt1	VM10RPTF	114	30:	
s3cavt2	VM10RPTF	116	31:	
s3cavt3	VM10RPTF	118	32:	
s3ctt1	VM10RPTF	121	33:	
s3ctt2	VM10RPTF	123	34:	
s3ctt3	VM10RPTF	125	35:	