Enterprise Systems Services and Support

VM/ESA Release 1.1 Performance Report

C. L. Morse
A. M. Shepherd

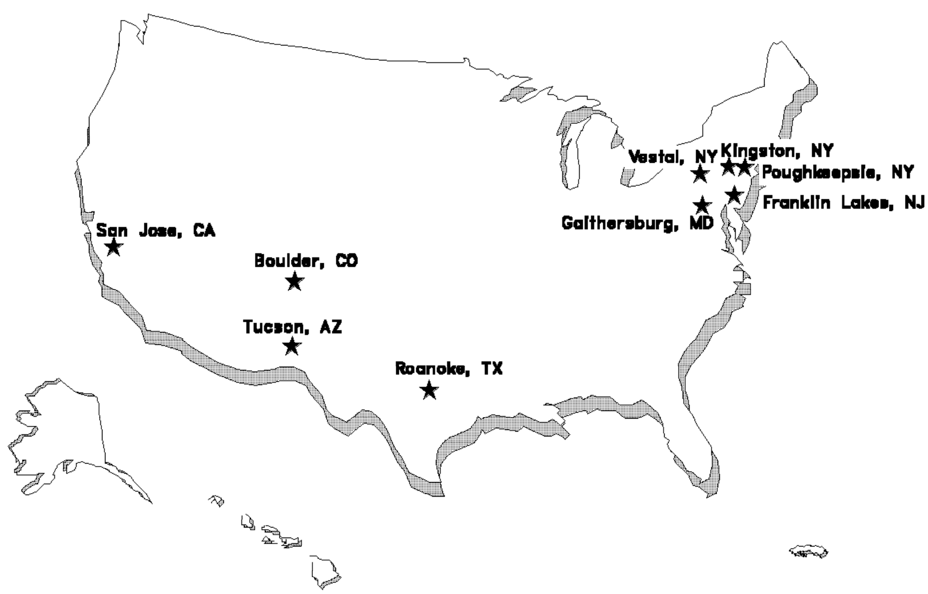

Washington Systems Center Technical Bulletin

GG66-3236-00

Enterprise Systems
Services and Support
Washington Systems Center
Technical Bulletin

VM/ESA Release 1.1 Performance Report

C. L. Morse
A. M. Shepherd

GG66-3236-00
April, 1992.

First Edition

This book applies to VM/ESA Release 1.1. (Program Number 5684-112).

Publications are not stocked at the address given below. Requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

A form for readers′ comment is provided at the back of this publication. If the form has been removed, comments may be addressed to: Chuck Morse, IBM Corporation, Washington Systems Center, 800 N. Frederick Avenue, Gaithersburg, MD 20879-3395.

Contents

Figures

Notices

The information contained in this document has not been submitted to any formal IBM test and is distributed on an “as is” basis **without any warranty either expressed or implied**. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Performance data contained in this document was determined in various controlled laboratory environments, and is for reference purposes only. **Customers should not adapt these performance numbers to their own environments as system performance standards.** The results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

This publication references specific APAR numbers which have an affect on performance. The APAR numbers included in this report may have pre-requisites, co-requisites, and/or fixes in error (PEs). The information included in this report is not a replacement for normal service research.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM licensed program in this publication is not intended to state or imply that only IBM′s program may be used. Any functionally equivalent program may be used instead.

Programming Interfaces

This publication is intended to help the customer understand the performance of VM/ESA Release 1.1 on various IBM processors. The information in this publication is not intended as the specification of any programming interfaces that are provided by VM/ESA Release 1.1. See the IBM Programming Announcement for VM/ESA Release 1.1 for more information about what publications are considered to be product documentation.

Trademarks

The following are trademarks of the IBM Corporation:

ECKD
ES/3090
ES/9000
IBM
VM/XA
VM/ESA
AIX
OS/2
PS/2
SAA
SQL/DS
System/370
System/390
VTAM
MVS/SP
OfficeVision
3090

Acknowledgements

The following personnel contributed to this report:

Endicott Programming Laboratory

M. S. Bidwell
W. J. Bitner
G. O. Blandy
C. T. Bradley
P. A. Bronson
R. A. Buck
P. A. Duncan
W. G. Ernsberger
G. S. Gasper
L. H. Guyer
G. A. Hine
R. R. Kaskan
S. P. Lyons
S. T. Marcotte
V. E. Meredith
L. L. Quinn
S. E. Shearer
A. M. Shepherd
D. P. Spencer
C. R. Terry
M. S. Thomas
J. L. Thrall
J. S. Tingley
A. M. Ward
M. G. Weber
R. C. Zinna

Kingston Programming Laboratory

S. P. Davis
C. L. Dewar
E. Digan
J. F. Eschmann
M. Ferrell
B. A. Hall
S. R. Hansell
R. A. Klein
F. D. Lewis
D. A. Vastola

Washington Systems Center

A. G. Coley
C. L. Morse

Southern Area Systems Center

G. S. Eheman

Abstract

The VM/ESA 1.1 Performance Report provides information on the performance of VM/ESA 1.1 running various workloads on the 9021, 9121, 9221, and 3090 processors.

The intent of this report is to provide performance and tuning information based on the results of the VM/ESA 1.1 performance test conducted jointly by the Endicott and Kingston programming laboratories.

Discussion centers on the performance effects of migrating from a previous VM release (usually VM/ESA 1.0) to VM/ESA 1.1, exploiting the new functions provided in VM/ESA 1.1, and using certain tuning options. In addition, some hardware processor and storage capacities are also included.

This report contains a General Observations section for those interested in a summary of the more significant findings and a Specific Measurements section for those readers interested in understanding the actual runs made and the complete analysis of the resulting data. In addition, this report contains performance experience from one ESP customer who migrated from VM/XA Release 2.1.

Part 1. Introduction

Organization of this Publication

This report provides performance information about various workloads with different hardware and software configurations running on VM/ESA 1.1 software. This report is divided into four parts:

1. Introduction

   This part gives an overview on the organization of this document and the hardware/software used for these measurements.

2. General Observations

   This part is a summary of the major conclusions reached as a result of the measurements performed for this report. It is recommended that the reader read this section first to get an overall idea of the performance effects of VM/ESA 1.1. If more information is desired on a topic, proceed to the same topic header under "Specific Measurements" to get detailed information.

3. Specific Measurements

   This part contains the measurement configuration used for the measurements that were run, the analysis of the results, and the table of performance data from which the conclusions were derived.

4. Appendixes

   This part contains additional performance data or additional information about the runs that are relevant to this report and its conclusions as well as the performance experience of one ESP customer and additional information on the CMS nucleus reduction. Appendix A and Appendix B contain additional performance data. Appendix C, D, and E all give additional information on the workloads, hardware configuration, and software configurations used for the measurements listed in this document. Appendix F discusses an ESP customer experience with performance when the customer migrated to VM/ESA Release 1.1. Appendix G contains the text of a Washington System Center (WSC) Flash describing how to improve performance by putting some CMS modules in logical shared segments.

The General Observations and Specific Measurements parts have been divided into four chapters:

- Migration/Regression

  This chapter analyzes the performance effects of migrating from the previous software level to the next release. Most of the measurements were done to determine what to expect when migrating from VM/ESA 1.0 (or earlier VM releases) to VM/ESA 1.1 for various environments (CMS intensive, OfficeVision, MVS Guest and VSE Guest). It also explores the effects of using storage above the 16M line, migrating to Shared File System (SFS), and migrating to the new software modes.

- Hardware Capacity

  This chapter analyzes the capacity of various 9021 and 9121 processors when running VM/ESA 1.1. In addition, several 9021-720 storage constrained runs are analyzed to show the effects of varying the amount of real and expanded storage on the performance of a CMS intensive workload.

- New Functional Enhancements

  This chapter analyzes new functions that are available with VM/ESA 1.1. Since there is nothing in a previous release to compare to the new function, the analysis compares the new function to the previous alternative or compares the option of using the new function against not using that function.

- Tuning Considerations

  This chapter analyzes the performance effects of using certain tuning options on VM/ESA 1.1. This report discusses (1) tuning several 9221 options when migrating to VM/ESA 1.1, (2) using XSTOR exclusively for minidisk cache on a 9121, (3) using the SET RESERVE command, and (4) turning off the OfficeVision message flags.

Hardware Used

The following processors were used for this report:

- 9021-720

  A pre-release equivalent of a 9021-720 was used for most of the measurements in this report. When an actual 9021-720 was available, some additional measurements (not shown in this report) were done which confirmed that the performance characteristics of the pre-release model are equivalent to the performance of the 9021-720. This pre-release processor was also used for the 9021-580, 9021-500, and 9021-340 measurements by varying off several processors of the 9021-720 to create the smaller model.

- 9021-580

  A pre-release equivalent of a 9021-580 was used for the 9021-580 CMS regression measurements only.

- 9121-480

  This processor was used for the 9121-480 and 9121-320 measurements. When running on a 9121-320, one processor was varied off to create the smaller model.

- 9221-170

  This processor was used for the 9221-170 measurements.

  The 9221 had 256M of real storage. Any real storage not defined for the specific measurement was configured as expanded storage and the excess was attached to an idle user.

- 3090-600J

  This processor was used for the 3090-600J and 3090-300J measurements. When running on a 3090-300J, three processors were varied off to create the smaller model.

Software Used

Unless otherwise noted, the early ship (ESP) level of VM/ESA 1.1 was used for the measurements in this report. When the GA level of VM/ESA 1.1 was available, several measurements (not shown in this report) were done to confirm that the performance characteristics of the VM/ESA 1.1 GA code were comparable to the VM/ESA 1.1 ESP level of code. These GA level measurements confirmed that the performance effects and conclusions listed in this report (based on the ESP level code) also apply to the GA code which our customers will be receiving when they order VM/ESA 1.1. All performance enhancements discussed in this report are part of the ESP level of code unless otherwise noted. Most of the known post-ESP items are mentioned in chapter 1, "Changes That Affect Performance" on page 7.

All previous releases of VM referred to in this report (ie: VM/ESA 1.0, VM/XA 2.1, etc) are at GA level. Unless otherwise noted, VM/ESA 1.0 refers to VM/ESA 1.0 ESA feature GA code. Thus, any service which was put into a previous VM release and forwarded to the VM/ESA 1.1 GA code, is not included in the previous release and can account for some of the difference between the previous release and the VM/ESA 1.1 release.

See the appropriate workload section in Appendix C, "Workloads" on page 273 for the other operating system and program product software levels.

Part 2. General Observations

1. Changes That Affect Performance

This chapter contains descriptions of various changes in VM/ESA 1.1 that affect performance. The majority of the changes are performance improvements, but some have potential for degradation. The objectives of this chapter are as follows:

- Provide a comprehensive list of the significant performance changes.
- Allow installations to determine for themselves whether their workloads would be affected by VM/ESA 1.1 performance changes.
- Describe new functions which applications could exploit to improve performance.

Throughout the rest of the report, various references are made to these changes. This further illustrates where the benefits occur.

Pending Page Release

Prior to VM/ESA 1.1, CMS issued DIAGNOSE X¢10¢ to release pages of storage. CP would reclaim host resources that had been required to support these pages. Due to the nature of the CMS Storage Manager and applications, these pages (virtual frames) are often requested for re-use in a short period of time. This behavior resulted in significant overhead in managing associated host resources.

In VM/ESA 1.1, CMS uses DIAGNOSE X¢214¢ in the management of page release. DIAGNOSE X¢214¢ provides functions to establish or cancel the pending release for a range of pages. This allows CP to delay or omit processing to reclaim host resources. A storage key option is also provided in the DIAGNOSE X¢214¢ functions. Unlike DIAGNOSE X¢10¢, DIAGNOSE X¢214¢ is not for general use. It is documented only in the CP Diagnosis Reference manual. The CMS SET RELPAGE OFF command is still respected.

This improvement addresses the majority of CMS environments. The use of DIAGNOSE X¢214¢ improves performance as follows:

- Reduces overhead in CP for page release.
- Decreases first time page faults.
- Reduces the number of PTLB (Purge Translation-lookaside Buffer) instructions.
- Allows CMS to avoid expensive SSKE instructions through use of the storage key option. This also eliminated the need for CMS to reference every page at IPL time, and thus eliminated the creation of associated page tables at that time.

RTM/ESA reports DIAGNOSE X¢214¢ on the DISPLAY PRIVOPS screen. DIAGNOSE X¢214¢ is included in the "Count of IBM supplied DIAGNOSE instructions" for each processor in the CP Monitor (Domain 0 Record 1).

CP Fast Dynamic Linkage

In VM/ESA 1.0, two forms of CP linkage exist: static and dynamic. Static linkage is efficient, but the implementation is more involved. Static linkage requires hardcoded, preallocated save areas. Static linkage is only used for a small number of frequently hit entry points. Dynamic linkage is much less efficient, but easy to program with. Analysis showed that VM/XA and VM/ESA 1.0 systems spend a considerable amount of time in HCPSVC which handles dynamic linkage. Linkage was chosen as an area to optimize. Fast dynamic linkage addresses this area. It is somewhere in between static and dynamic linkage, being relatively efficient and easy to work with.

There are some restrictions associated with fast dynamic linkage. These are:

- The called module must be resident.
- Fast dynamic linkage can not be used by a multiprocessor module to call a non-multiprocessor module (master only module).

The following factors lead to the performance improvement:

- Code is in line and therefore avoids a costly call to HCPSVC. Note that there is an option to use HCPSVC.
- No check is made to determine if the module is pageable at execution.
- No check is made to determine if the module is master only at execution.

Also, the cross processor return queue for save blocks is not used with fast dynamic linkage. This is not a problem since any temporary imbalance is corrected by normal dynamic linkage use.

Approximately 40 entry points from a total of 20 key modules were converted to use fast dynamic linkage in VM/ESA 1.1. These can be determined by looking at the HCPMDLAT MACRO.

In Line Page Table Invalidation Lock

VM/ESA 1.0 introduced the ability to page Page Management Blocks (PGMBKs). At that time, the page table invalidation lock (VMDPTIL) was made a formal lock managed by the HCPLCK module. Due to resources required to manage a formal lock and the frequency at which the lock was obtained/released, it was shown to be a potential area for significant improvement. This item implemented efficient in line macros to handle the most frequent scenarios. The in line macros handle the following scenarios:

- Wish to acquire lock and it is available.
- Wish to release lock and there are no queued requests.
- Wish to swap lock and there are no additional owners or queued requests. Swap is between *shared* and *exclusive* modes.

The original change in VM/ESA 1.0 to a formal lock also made holding the VMDPTIL a *critical process*. For each virtual machine, a count of critical processes is kept. These are meant to represent locks or resources held that are critical to system performance. If the count of critical processes is non-zero, CP temporarily gives the virtual machine special treatment to avoid a virtual machine′s delay from impacting other users. This is known as a *lock shot*. It was determined that the VMDPTIL was not a critical process so the overhead of maintaining the critical process structure could be avoided.

CCW Fast Path

A major weakness of VM/XA and VM/ESA 1.0 compared to VM/SP HPO is the additional CPU requirements associated with I/O. The majority of the increase is in CCW translation. Some areas of this are addressed by VM/ESA 1.1. Improvements to CCW processing for virtual machine DASD I/O were made, involving both translation (from virtual to real) and untranslation (from real to virtual).

Performance efficiencies were gained by adding a fast path for CCW translations. This change is referred to as CCW fast path. Benefits result from recognizing and optimizing for typical channel program structures. The code is optimized for successful I/O completion. Processing associated with the handling of error conditions is only performed when the error conditions occur. The processing required for untranslation is now a small and fixed amount. It is assumed the I/O will complete successfully (which it does the majority of the time). The longer the channel program, the greater is the improvement. This is due to the low and fixed costs for untranslation.

The CCW fast path function is not included in the VM/ESA 1.1 base GA code, but was added later by APAR VM51012. Since it was integrated late in the cycle, only a few of the measurements included in this report include CCW fast path. For CMS file I/O, only the DIAGNOSE X¢A8¢ DASD I/Os were affected by this support. The FS7B workloads do relatively little DASD I/O via DIAGNOSE X¢A8¢ and therefore see only a small benefit from this change. I/O intensive commands like DDR and FORMAT showed significant improvement. In general, CCW fast path does not close the gap between VM/HPO and VM/ESA CCW translation costs for CMS intensive environments. CCW fast path greatly benefits many V=V guest operating system environments when compared to VM/XA or VM/ESA 1.0 ESA feature.

MDC Spin Lock Fix

In environments with large amounts of expanded storage being used heavily by minidisk caching, the potential existed for sporadic periods of very high lock spin time. Minidisk cache management requires certain system locks when reorganizing its hash table structure. This caused high lock spin time spikes seen with RTM as %SP on the DISPLAY CPLOG screen. Depending on when RTM was last reset, the problem is not always obvious by looking at the average. The severity of the problem is proportional to the number of processors on the system, the degree of MDC activity, and the size of the minidisk cache.

Changes went into the minidisk cache management processing to eliminate the cause of this problem. These are in the base of VM/ESA 1.1 and were added to VM/XA 2.0, VM/XA 2.1, and VM/ESA 1.0 via APARs. The APARs are as follows:

- VM/XA 2.0 - VM44286 + PE VM44894
- VM/XA 2.1 - VM44286 + PE VM46595
- VM/ESA 1.0 - VM45731 + PE VM45398

Please note that the above are base and fix APARs, but normal service research should be done for co-requisites, pre-requisites, and fixes in error.

IUCV Improvements

This includes both IUCV and APPC/VM. With the growing reliance on server virtual machines, the need for efficient communication functions grows. IUCV and APPC/VM were shown to be more expensive in VM/ESA than in the 370 based CPs. These factors led to a focus on improving IUCV performance.

Several changes led to the improved performance:

- Storage management for MSGBKs and IUSBKs was made more efficient. Semi-permanent control blocks and stack management were exploited.
- Several high frequency entry points were converted to fast dynamic linkage (previously mentioned).
- The processing of external interrupts was optimized.

These changes resulted in improved performance for base IUCV and APPC/VM functions. The improvement is on a per IUCV basis, not a per byte transferred basis. Therefore, the percentage improvement was greater for smaller size transactions. This is shown by results from a single thread APPC/VM benchmark:

- 512 Byte Transactions
  - Total CPU down 20% per transaction
  - Transaction rate up 24%
- 100K Byte Transactions
  - Total CPU down 3% per transaction
  - Transaction rate up 2.5%

Both VTAM and SFS use IUCV or APPC/VM and benefitted from this item.

GCS IPOLL

Changes in GCS and CP introduced the IPOLL function, which GCS can use to poll for pending replies and messages. By allowing GCS to retrieve up to 102 “interrupts” on each IPOLL, there is a potential for significant reduction (e.g. 100 to 1) in the number of interruptions to GCS applications when they are flooded with incoming IUCVs. AVS and VSCS are examples of such applications. The performance improvement is proportional to the IUCV interrupt activity. Therefore, large systems tend to see more benefit than small systems. See “GCS IPOLL Option” on page 40 for general observations and “GCS IPOLL Option” on page 207 for measurement details.

Inter-System Facility for Communication

The Inter-System Facility for Communications (ISFC) is a new component in VM/ESA 1.1 systems that provides high-speed connectivity to groups of LAN attached workstations running VM PWSCS. This function is implemented directly in CP. Superior performance is achieved with this function imbedded in CP, as opposed to a server virtual machine implementation. See “Inter-System Facility for Communication (ISFC)” on page 41 and “Inter-System Facility for Communication (ISFC)” on page 213 for additional information.

Block Paging Improvements

The efficiency of block paging is decreased when the optimal blocking factor can not be used. A number of things that used to prevent this or "break" the block were removed:

- Blocks not broken on segment faults.
- Blocks not broken when available list is empty.
- Blocking up to 64 pages.
- Blocking as large as the virtual machine specifies via REFPAGE CP macro. REFPAGE is new and involves VM Data Spaces. It deals with an application giving CP hints about the virtual machine page references. See "Page Reference Pattern Function" on page 13 for more details.

Default DSPSLICE

When a VM/ESA system is IPLed, CP initialization logic uses a timing loop to determine the default dispatching minor time slice. This is an attempt to determine an appropriate value for the speed of the processor VM is running on. Analysis and experimentation showed that the default calculation resulted in a less than optimal value for several high end processors. Further analysis showed that a floor of 5 milliseconds provided improved ITR for the affected processors.

In VM/ESA 1.1, the initialization logic is the same for the calculation of the default dispatch slice, except an additional check is added. If the computed default value is less than 5 milliseconds, it is changed to be 5 milliseconds. The value can still be changed via the SET SRM DSPSLICE. The range of acceptable values stays the same (1 to 100 milliseconds). The current setting can be determined with the QUERY SRM DSPSLICE command. An APAR for this change also exists for VM/ESA 1.0 (VM48108).

The benefit associated with this change is dependent on the workload and processor. Results showed an ITR increase in range of 0.3% to 3.3%. The ITR improvements were a result of less CP resource. This occurs when fewer time slices are required per user transaction, and thus less CP dispatcher processing. For example, if the old value was 2 milliseconds and the majority of transactions required 3 milliseconds to complete, the new value (5 milliseconds) would allow these transactions to complete in a single dispatch time slice.

However, there are some other considerations to this change. Most measurements also showed improved response time corresponding to ITR change. The exceptions were systems with remote users connected via VTAM and VSCS. In these environments, it is believed that the increased dispatch time slice made VTAM less responsive, which resulted in response time staying the same or becoming somewhat worse.

In addition to the response time factors, the following need to be kept in mind:

- Impact to explicit tuning where the dispatching minor time slice is a factor. For example, the SET SRM IABIAS command parameters are related to the minor time slice.
- Relationship to master processor constraints since a virtual machine may run on the master processor for longer periods of time before CP gains control.

- Impact to run away or CPU intensive users since these virtual machines may run for longer periods of time before CP gains control.

Some of the above items relate back to the CP dispatcher getting control less frequently and therefore being less responsive in taking corrective action.

VM Data Spaces

VM Data spaces provide increased storage addressability and therefore can move the burden of I/O from an application to CP. The use of storage as a buffer is an old performance concept. An application may still be responsible for initially loading data into storage, but CP paging handles it after that.

The use of VM Data Spaces also extends the concept of sharing data. This has two chief advantages:

1. It reduces storage requirements. One copy can be shared among virtual machines instead of a copy per virtual machine.
2. It reduces the need to transfer data by IUCV, APPC/VM, or some other communication vehicle.

VM/ESA 1.1 introduces the new virtual machine mode of XC for exploitation of VM Data Spaces. For non-XC mode virtual machines, DIAGNOSE X¢248¢ (Copy to Primary function) can be used to move data from data space to primary address space. Callable Services Library (CSL) provides an interface with high level language support.

See the list of references at the end of this document for sources of additional information.

Minidisk Mapping: Minidisk Mapping extends the concept of applications using storage and letting CP handle the real I/O. This new function provides a means of associating CP minidisk data with data spaces. One or more minidisks can be mapped to one or more data spaces. An application retrieves the data simply by referencing the corresponding storage location in the data space. The real I/O is handled by the CP paging subsystem, which provides efficiencies over virtual machine I/O.

Some initial setup work is required to establish the mapping rules. This is managed by MAPMDISK, a CP macro. Since virtual storage is volatile, management for integrity must be considered. The SAVE function provides a means of forcing (committing) the data to the nonvolatile DASD where the minidisks reside.

Asynchronous Page Faults: There is growth in the number and importance of server virtual machines on VM. One problem associated with server virtual machines is the impact of serialization on other dependent virtual machines. The processing associated with page fault resolution serializes the virtual machine. When a server virtual machine is serialized, so are the dependent end user virtual machines. Asynchronous Page Fault capability is introduced to help these scenarios.

Asynchronous Page Fault allows a virtual machine to continue processing other work (a different task), while CP handles the page fault. The implementation applies only to page faults of data space pages. CP will provide an interrupt when the page fault is complete. At that time, the server application can finish processing the original task associated with the page fault.

The server application requires logic to work in this environment. This includes:

- A structure that lends itself to tasking or work units.
- Selection of asynchronous page fault function on a data space by data space basis. This occurs when adding the data space to the access list.
- Using the CP macro PFAULT to establish a token for handshaking.
- Support to handle the associated interrupts.

Some will point out that this concept is not new to VM. The PAGEX support is based on the same idea. There are two significant differences between the two:

1. PAGEX deals with the primary address space while the Asynchronous Page Fault support is limited to VM Data Spaces.
2. Asynchronous Page Fault was designed with server virtual machines in mind. The handshaking interface with CP is easy to work with and lends itself nicely to server applications.

Page Reference Pattern Function: The most efficient page fault is the one that does not occur. The Page Reference Pattern Function addresses this. CP logic attempts to maintain the appropriate working set for virtual machines, but it can not predict the future. CP can only guess as to which pages a virtual machine will reference in the future. Page Reference Pattern provides a method for the application to give hints to CP as to which pages will be referenced in the very near future.

This is accomplished with the REFPAGE macro. The pages need not be contiguous. The REFPAGE macro can be used for the primary address space or data spaces. In all cases, the virtual machine must be running in XC mode for Page Reference Pattern. The effect is very short term. After the REFPAGE is issued, all non-resident pages will be brought in as a block set at the time of the first page fault for the pattern. After that, a new REFPAGE would be needed. Misuse of this function could result in performance problems instead of benefits.

SFS Performance Improvements

SFS Exploitation of VM Data Spaces: VM Data Space usage in general has been discussed. SFS exploits many of the data space aspects to provide improved performance for DIRCONTROL directories containing read-mostly data. SFS exploitation improves performance by avoiding file pool server requests and by the sharing of data within the data space. The data space is owned and maintained by the file pool server.

The actual contents of the data space are as follows:

- The shareable part of the Active Disk Table (ADT) control block.
- File Status Table (FST) control blocks. Previously, only EDF minidisks could share FSTs. (EDF accomplishes this by using saved segments created with the SAVEFD command.)
- File data blocks. The minidisk mapping functions are used to maintain these.

Additional information can be found in "Exploitation by Shared File System" on page 37 (general observations) and "VM Data Spaces: Exploitation by Shared File System" on page 161 (measurement details).

SFS Checkpoint Improvement: Checkpoint processing, a normal part of managing the SFS logs, causes serialization of the file pool server. As mentioned earlier, serialization is bad for servers and affects user response time.

This improvement doubled the interval between checkpoints. It is now done every 100 filled log pages as opposed to 50 in VM/ESA 1.0. This improves response time. However, since checkpoints are relatively infrequent, there is no significant reduction in I/Os or processor usage.

SFS Asynchronous File Functions: Earlier versions of SFS were able to optionally perform some functions asynchronously. VM/ESA 1.1 extends this to several key functions including Open, Close, Read, and Write. These are the key functions of any file system. The new asynchronous functions allow applications to exploit parallelism or run as server machines without being serialized by their outstanding SFS requests.

CSL Loaded above the 16M Line

CSL is the Callable Services Library and provides many routines callable from high level languages. CSL is loaded at IPL time by the SYSPROF EXEC via the RTNLOAD command. It is located in the user′s virtual storage, occupying more than 350 Kb in VM/ESA 1.0. This led to an increase in “virtual storage exhausted” messages. The capability exists to use a saved segment and add SEGMENT LOAD to SYSPROF prior to the RTNLOAD command, allowing a shared copy of CSL to be used. However, prior to VM/ESA 1.1, CSL only ran below the 16M line and many sites do not have much room in that area. In VM/ESA 1.1, CSL can run above the 16M line. This change is available in VM/ESA 1.0 via APARs VM44717 and VM47566. Please note that the above are base and fix APARs, but normal service research should be done for co-requisites, pre-requisites, and fixes in error.

CMS Nucleus Growth Relief

A large amount of new function has gone into CMS in the past few releases. That new function has also led to a growth of the CMS nucleus. In fact, the rate of growth has also increased. In VM/ESA 1.0, the growth resulted in some installations not having enough room for the S and Y Stats (saved FSTs for the S and Y disks). This can lead to performance problems. In VM/ESA 1.1, continued growth would have lead to CMS requiring an additional segment in the already crowded area below the 16M line. Two key changes were made to address this. The first was in the management of the CMS message repository. The second was moving some CMS code from the nucleus to the S-disk.

Analysis showed that the greatest single source of growth was the CMS message repository. There were messages associated with all the new function that had been going into CMS. In VM/ESA 1.1, the repository now starts at the 16M line. Access for XA and XC mode is straightforward. In 370 mode, a version of DIAGNOSE X¢248¢ (copy-to-primary) is used. For all modes, the management of messages was enhanced to provide true caching. In the past, some key messages were cached by hard coding them to avoid message repository processing.

To further reduce the size of the CMS nucleus some commands were removed from the nucleus and placed as modules on the S-disk. (This change was introduced in VM/ESA 1.0 by APAR VM49762 and is included in the VM/ESA 1.1 base.) A total of sixteen modules were moved. While an attempt was made to ensure that performance sensitive modules were not removed from the nucleus, some environments may require the use of a subset of the commands. Invoking a module residing on the S-Disk results in it being loaded into the end user′s virtual storage as a nucleus extension. This storage is not shared and therefore can cause performance degradation in storage constrained environments due to the increase in user working set size and system paging.

One of the steps that can be taken to offset this effect is to place some or all of these modules into a logical shared segment, thus allowing all users of these modules to share a single copy. The steps necessary to do this were documented in a Washington System Center flash in February of 1992 and are also included in Appendix G.

Disconnect/Reconnect Handshaking

Users expect CMS to be able to handle the scenario where they disconnect from one terminal and reconnect on a different size terminal. CMS is expected to adjust to the new screen size. In the past, CMS was constantly issuing DIAGNOSEs for terminal characteristics in order to accomplish this. This is costly in terms of CP processing associated with the DIAGNOSE handling and the SIE breaks that were caused. Now, CP and CMS shake hands via a new DIAGNOSE (X¢264¢). DIAGNOSE X¢264¢ is not for general use. During CMS IPL processing, CMS will issue the DIAGNOSE to inform CP of a communication area. A flag is established that CP will use to indicate that CMS needs to redetermine screen characteristics. CMS merely checks the flag in virtual storage instead of continuously issuing DIAGNOSEs.

XA Mode CMS Improvements

A performance goal for VM is to narrow the gap between 370 mode and XA mode virtual machine performance. Changes were made in VM/ESA 1.1 to help close the gap. Several changes were made to eliminate expensive privileged instructions in XA mode paths. See “Software Mode Comparisons” on page 28 (general observations) and “Software Mode Comparisons” on page 119 (measurement details) for additional information.

CRR LUWID Pooling

This was really a positive side effect of a general change for SFS. It was determined that there were some obscure situations in which we could get into an undetectable deadlock with SFS. These scenarios were quite complex. In order to resolve these, the SFS file pool server needed to be passed a global LUWID (logical unit of work identifier) from the end user for each file pool request. The CRR (Coordinated Resource Recovery) server manages these LUWIDs. In the past, CMS code in the end user virtual machine would request a LUWID from the CRR server and a single LUWID was passed back.

The additional CRR server requests would have been a performance problem. Therefore, LUWID processing was changed to have the CRR Recovery server return multiple LUWIDs (255). The impact to normal SFS regression performance is negligible. However, for CRR exploitation cases this results in a significant performance improvement. A CRR server request to get an LUWID is needed only once every 255 commits instead of for every commit.

CMS Pipelines

CMS Pipelines was previously available as a PRPQ, but is part of base VM/ESA 1.1 with APAR VM47212. Actually, the code is in the GA level of VM/ESA 1.1 and the APAR makes documentation, help files, and support available. CMS Pipelines provides function to use output from one program as input to another program. This redirection can be repeated across several programs. In addition, CMS Pipelines provides a set of functions that can be used to manipulate data between programs. In some cases, performance improvements can be gained by using CMS Pipelines instead of combinations of REXX, EXECIO, and XEDIT. In general, the performance gains are associated with using CMS Pipelines for additional data manipulation, not just redirection. See "CMS Pipelines" on page 39 and "CMS Pipelines" on page 196 for more details.

DASD Fast Write

DASD fast write is a 3990-3 feature which decreases write response time by immediately returning channel end and device end when a write hit occurs. The 3990 controller does the actual write asynchronously when the device is available. The data is saved in nonvolatile storage (NVS), thereby eliminating the possibility of data loss even if a power failure occurs.

Environments with heavy write I/O activity can benefit from exploiting DASD fast write. This is especially true when the I/O is performed synchronously by a server application. See the following sections for more details: "3990-3 DASD Fast Write" on page 38 and "3990-3 DASD Fast Write" on page 180.

See the list of references at the end of this document for sources of additional information.

2. Migration/Regression

CMS Intensive Migration from VM/ESA 1.0

The CMS intensive regression measurements were made using the minidisk and 35% SFS workloads on selected 9021, 9121, and 9221 processors.

The performance of VM/ESA 1.1 showed improvement over VM/ESA 1.0. All measured environments showed an increased internal throughput rate (ITR), lower response times, and reduced processor utilization. The amount of improvement is a function of the processor configuration and the level of SFS usage. The ITR improvement ranged from 3.2% to 6.9% while the external response time decrease ranged from 0.04 seconds (7.6%) to 0.29 seconds (35.5%). These improvements are illustrated in Figure 1 on page 18 and Figure 2 on page 18.

There were several performance changes that went into VM/ESA 1.1 and which are discussed in chapter 1, “Changes That Affect Performance” on page 7. There was some growth in real storage requirements but this was outweighed by several performance improvements. The following improvements provided the most benefit to the CMS intensive environment:

- **CP Fast Dynamic Linkage.**
- **IUCV Improvements.** The IUCV improvements will show more benefit in environments having higher APPC/VM usage. For instance, measurements with external VSCS virtual machines showed greater improvements in ITR than those with internal VSCSs. Since SFS uses APPC/VM for communications with user machines and uses the *BLOCKIO interface for I/O, the IUCV improvements help the SFS environment more than the minidisk case, resulting in additional throughput and response time benefits.
- **Pending Page Release.** The Pending Page Release improvement decreases the number of SSKE and PTLB instructions which are more expensive when there are multiple processors due to processor signaling. Thus, this change shows greater improvements in large multiprocessor environments.
- **MDC Spin Lock Fix.** The MDC Spin Lock Fix helps environments with multiprocessors and large amounts of minidisk cache (MDC). It provides no benefit on uniprocessors such as the 9121-320.
- **XA-Mode CMS Improvements.**
- **SFS Checkpoint Improvement.** The SFS Checkpoint Improvement accounts for most of the additional improvement in response time relative to the minidisk environment.

The improved ITRs for VM/ESA 1.1 are due to decreased CP overhead which was brought about by these changes. These also influence internal and external response times.

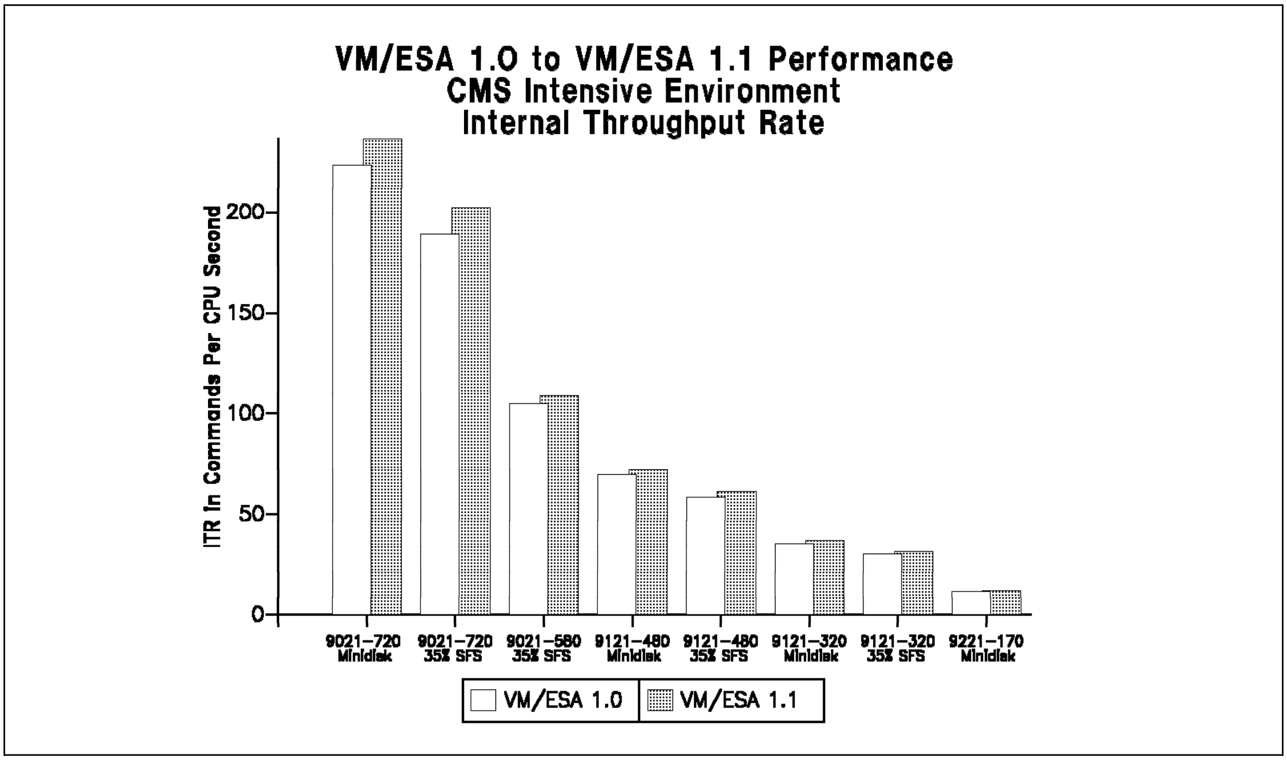

Figure 1. Internal Throughput Rate for the FS7B Workload

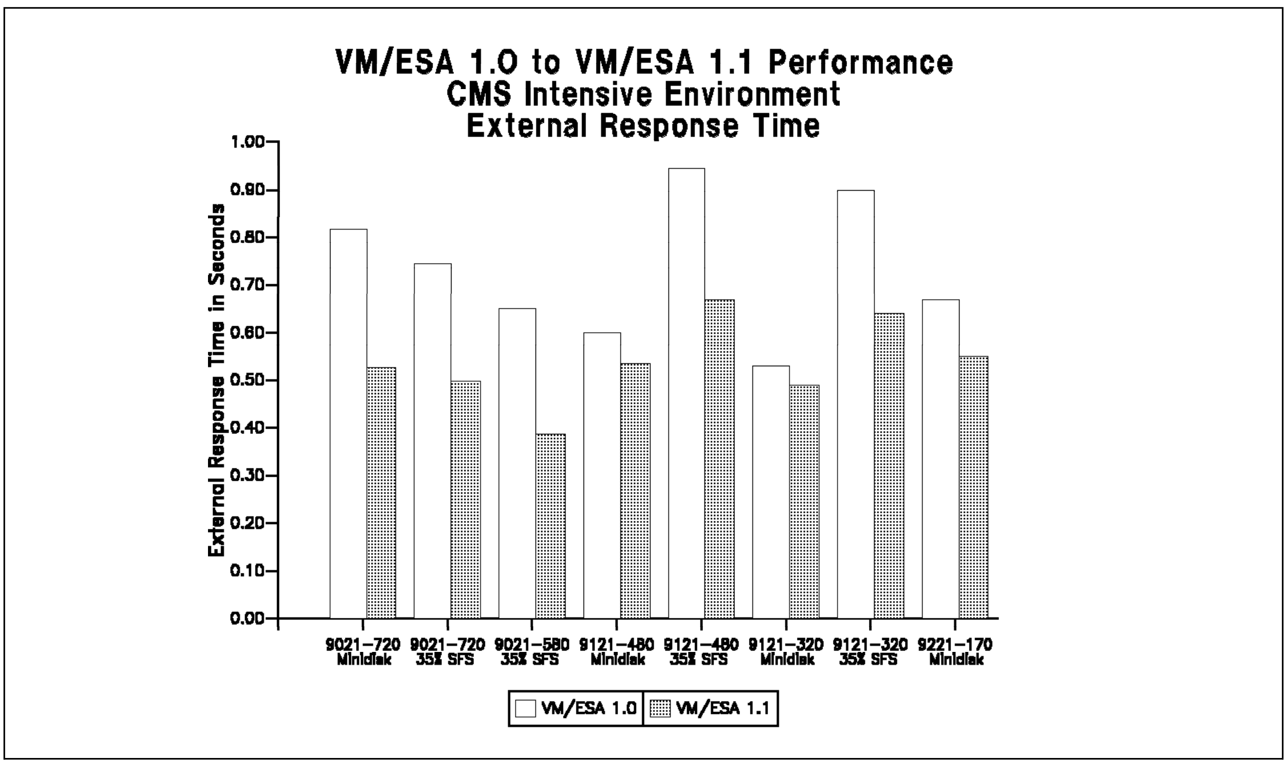

Figure 2. External Response Times for the FS7B Workload

CMS Intensive Migration from VM/ESA 1.0 370 Feature on 9221

The following table summarizes the 9221-170 migration measurements from VM/ESA 1.0 370 Feature to VM/ESA 1.0 and to VM/ESA 1.1. The measurements are made using the minidisk workload, 35% SFS workload, and maximum SFS with VM Data Spaces workload. Most of the measurements that appear in the same row or the same column are compared in the appropriate sections of this report.

Explanation of columns:

- MEASUREMENT DESCRIPTION
  - Users
    - minidisk (EDF) or SFS measurement
    - the number of users in the measurement (280 or 240)
    - tuned measurement (yes) or untuned measurement (no)
  - Storage
    - real storage (first number)
    - expanded storage (second number)
- TABLE ENTRY DESCRIPTION

  The items in this column provide a description of the information in each row of the results.
- TABLE OF RESULTS BY SYSTEM ENVIRONMENT

  Each column represents a unique combination of VM system level and user mode.

Explanation of rows:

- VM System

  level of VM
- User mode

  user virtual machine mode
- Run ID

  measurement identification
- ITR Ratio

  internal throughput rate ratio relative to the VM/ESA 1.0 370 Feature measurement being 1.000
- AVG LAST (T)

  external response time

Measurement Description		Table Entry Description	Table of Results by System Environment					
		VM System	ESA 1.0 370 Feature	ESA 1.0 ESA Feature	ESA 1.1 ESA Feature			
Users	Storage	User Mode	370	370	370	XA	XC	XC with Dataspaces
EDF/280/NO	64/0M	Run Id	H17R0281		H14R0283			
		ITR Ratio	1.000		0.863			
		AVG LAST(T)	0.800		0.890			
EDF/280/YES	48/16M	Run Id		H13R0280	H14R0287	H14R0289		
		ITR Ratio		0.906	0.945	0.919		
		AVG LAST(T)		0.670	0.550	0.630		
EDF/280/YES	240/16M	Run Id			H14R0286			
		ITR Ratio			1.007			
		AVG LAST(T)			0.430			
SFS/240/NO	64/0M	Run Id	H17F0241					
		ITR Ratio	0.872					
		AVG LAST(T)	0.787					
SFS/240/YES	48/16M	Run Id			H14F0241		H14F0242	H14M0241
		ITR Ratio			0.806		0.789	0.791
		AVG LAST(T)			0.690		0.710	0.660

Table 1. 9221-170 migration

Direct migrations from VM/ESA 1.0 370 Feature to VM/ESA 1.1 show that VM/ESA 1.1 has lower internal throughput and higher response times. For details, see section “9221-170 / Minidisk” on page 86. The internal throughput decrease is worse than the comparison between HPO 5 and VM/ESA 1.0 on a 3090-200J that was published in *VM/ESA Release 1.0 Performance Report*, ZZ05-0469.[1] A portion of this decrease is caused by the ESA mode implementation and a portion is caused by efficiency of VM/ESA 1.0 370 Feature on a uniprocessor versus a two-way processor. By implementing performance tuning options and adding more real storage, significant improvements in internal throughput and response times were made. For details, see section “Recommended 9221 Tuning” on page 223.

Note: 9221 processors configured with integrated I/O controllers, running in ESA/390 mode, can only use 128 MB of main storage. If the installed processor storage on these machines is greater than 128 MB, the remaining storage may be used for expanded storage. All of the runs shown above used only channel-attached devices, making the use of more than 128 MB of main storage possible.

1 This document is classified as IBM Internal Use Only. Contact you IBM representative for access to the information contained in this publication.

Migrating from VM/ESA 1.0 ESA Feature to VM/ESA 1.1 shows internal throughput improving by 4.3%. This is due largely to CP Fast Dynamic Linkage and IUCV improvements. For details, see "9221-170 / Minidisk" on page 82. Since the 9221-170 is a uniprocessor, the MDC Spin Lock Fix provides no internal throughput improvements. As explained in "9021-580 / 35% SFS" on page 62, the Pending Page Release enhancement has less effect on systems with fewer processors. External response time improved 17.9%.

Comparing the 9221-170 to other ES/9000 systems on VM/ESA 1.1, the internal throughput delta is :

- larger between 370 and XA mode virtual machines (see section "9221-170 / Minidisk" on page 123).
- consistent between minidisk and SFS (see section "Minidisk to Shared File System" on page 26).
- consistent between minidisk and VM Data Spaces (see section "VM Data Spaces" on page 37).

CMS Intensive Migration from Currently Supported Releases

A large body of performance information for the CMS intensive environment has been collected over the last several releases of VM. This section summarizes the internal throughput rate (ITR) data from those measurements in order to show the approximate changes in processing capacity, for CMS intensive workloads, that will tend to occur when migrating from one VM release to another. As such, this section can serve as one source of migration planning information.

The performance relationships shown in this section are limited to the CMS intensive environment. Other types of VM usage may show different relationships. Further, any one measure such as ITR cannot provide a complete picture of the performance differences between VM releases. The VM performance reports from which the ITR ratios (ITRRs) have been extracted can serve as a good source of additional performance information. These reports are listed at the end of this section.

Table 2 summarizes the ITR relationships we have observed for the CMS intensive environment for a number of VM release to release transitions:

From	To	Case	ITRR	ITRR Derivation	Notes
VM/SP 5	VM/SP 6		0.82	R5=avg(1,2)	1
	VM/ESA 1.0 (370)		0.88	R5*R6	1,5
	VM/ESA 1.0 (ESA)	4381	0.74	R5*R6*R13A	1,2,5
	VM/ESA 1.1	4381	0.77	R5*R6*R13A*R1E	1,2,5
	VM/ESA 1.0 (ESA)	9221	0.80	R5*R6*R13B	1,5,6
	VM/ESA 1.1	9221	0.83	R5*R6*R13C	1,5,6
VM/SP 6	VM/ESA 1.0 (370)		1.07	R6=avg(3,4)	5
	VM/ESA 1.0 (ESA)	4381	0.90	R6*R13A	2,5
	VM/ESA 1.1	4381	0.93	R6*R13A*R1E	2,5
	VM/ESA 1.0 (ESA)	9221	0.97	R6*R13B	5,6
	VM/ESA 1.1	9221	1.02	R6*R13C	5,6
VM/ESA 1.0 (370)	VM/ESA 1.0 (ESA)	4381	0.84	R13A=(note 2)	2,5
	VM/ESA 1.1	4381	0.87	R13A*R1E	2,5
	VM/ESA 1.0 (ESA)	9221	0.91	R13B=13	5,6
	VM/ESA 1.1	9221	0.95	R13C=14	5,6
VM/SP HPO 5	VM/ESA 1.0 (ESA)	MP, 4381	0.97	RH=8	3,4,5
	VM/ESA 1.1	MP, 4381	1.01	RH*R1E	3,4,5
VM/XA 2.0	VM/XA 2.1		1.02	R20=5	
	VM/ESA 1.0 (ESA)		1.06	R20*R21	
	VM/ESA 1.1		1.10	R20*R21*R1E	
VM/XA 2.1	VM/ESA 1.0 (ESA)		1.04	R21=avg(6,7)	
	VM/ESA 1.1		1.08	R21*R1E	
VM/ESA 1.0 (ESA)	VM/ESA 1.1		1.04	R1E=avg(9-12)	

Table 2. Approximate VM Relative Capacity: CMS Intensive Environment

Explanation of columns:

Case - The set of conditions for which the stated ITRR applies. When not specified, no large variations in ITRR have been found among the cases that have been measured. However, smaller variations are typically seen. These ITRR variations are shown in "Derivation and Supporting Data" below.

ITRR - The "to" ITR divided by the "from" ITR. A number greater than 1.00 indicates an improvement in processor capacity.

ITRR Derivation - Shows how the ITRR was derived. See "Derivation and Supporting Data" below for discussion.

Notes - The numbers shown refer to the following notes:

1. The VM/SP 5 system is assumed to include the performance SPE which adds segment protect and 4K key support. Other measurements have shown that VM/SP 5 ITR is 4% to 6% lower without this SPE.

2. The VM/ESA 1.0 (370) to VM/ESA 1.0 (ESA) comparison is based upon measurements done on a 4381-91E using the PD3 and HT4 CMS intensive workloads. Additional measurements with FS7B show similar results.

3. The VM/SP HPO 5 to VM/ESA 1.0 (ESA) comparison was done with reduced think time in order to avoid a 16M line real storage constraint in the HPO case. In cases where the base HPO 5 system is 16M line con-

strained, migration to VM/ESA will realize additional performance benefits due to the elimination of this constraint.

4. This comparison is based on measurements done on a 3090-200J multiprocessor. Since MP support is standard with VM/ESA but can be generated out for HPO, a less favorable ITR ratio can be expected for uniprocessors.

   The ESA-capable 4381 models provide less processing capacity when run in ESA mode as compared to 370 mode. Therefore, a less favorable ITR ratio can be expected when migrating a 4381 configuration from VM/SP HPO 5 to VM/ESA 1.0 (ESA) or VM/ESA 1.1.

5. The target VM system supports a larger real memory size than the stated migration source. This potential benefit is not reflected in the stated ITR ratios. Migrations from memory-constrained environments will tend to see additional ITRR and other performance benefits when the target configuration has additional real storage.

6. These results apply to the case where the following recommended tuning is done for the target system:

   Configure 16M as expanded storage for minidisk caching.
   Set the VTAM delay to 0.2 msec. (Default is no delay.)
   Preload the shared segments.
   Set DSPSLICE to three times the default.

   The purpose of this tuning is to optimize VM/ESA for use on ESA mode 9221 processors. If this tuning is not done, lower ITR ratios will be experienced. For example, for the FS7B0R CMS intensive workload, going from VM/ESA 1.0 (370 Feature) to VM/ESA 1.1 resulted in an ITRR of 0.95 with the above tuning and an ITRR of 0.86 without it. See "Recommended 9221 Tuning" on page 223 for further discussion of these tuning recommendations.

Bear in mind that this table only shows relative performance in terms of ITR ratios (processor capacity). It does not directly show how any two VM releases would compare in terms of response time. An improved ITR tends to result in better response times and vice versa. However, exceptions can occur. Also, the effect of ITRR on response times can, in an actual migration, be outweighed by other factors (such as hardware and workload) that have changed at the same time.

This table represents CMS intensive performance for the case where all files are on minidisks. The release-to-release ITR ratios for the case of SFS usage are very similar to the ones shown here. SFS release-to-release measurement results are provided in references 1 and 2 (listed at the end of this section).

These VM ITRR estimates can be used in conjunction with the appropriate hardware ITRR figures in order to estimate the overall performance change that would result from a migration that involves both a hardware upgrade and an upleveling of VM. For example, suppose that the new processor has an ITRR of 1.30 for CMS intensive workloads relative to the current system and further suppose that the migration also includes an upgrade of VM from VM/XA 2.1 to VM/ESA 1.1. From the above table, the estimated ITRR for migrating from VM/XA 2.1 to VM/ESA 1.1 is 1.08. Therefore, you would estimate the overall increase in system capacity as 1.30*1.08 = 1.40.

Derivation and Supporting Data

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

The derivation column in Table 2 on page 22 shows how the stated ITR ratio was calculated. For example, the ITRR of 1.08 for going from VM/XA 2.1 to VM/ESA 1.1 was calculated by multiplying the average ITRR for going from VM/XA 2.1 to VM/ESA 1.0 ESA Feature (R21) by the average ITRR for going from VM/ESA 1.0 ESA Feature to VM/ESA 1.1 (R1E): 1.08 = 1.04*1.04. R21 was calculated by averaging the ITRRs for VM measurement pairs 6 and 7 (see Table 3 on page 25). Likewise, R1E was calculated by averaging the ITRRs for VM measurement pairs 9 through 12.

Any given measurement pair represents two measurements where the only difference is the VM release level. As such, all the performance results obtained for one of the measurements in the pair can validly be compared to the corresponding results for the other measurement.

By contrast, there are often substantial environmental differences between measurement pairs. Factors such as number of users, workload, processor model, and I/O configuration will often be different. This greatly limits the kinds of valid inferences that can be drawn when trying to compare data across two or more measurement pairs. For example, response times are very sensitive to a number of specific environmental factors and therefore should only be compared within a set of controlled, comparable measurements.

For this reason, the above table only covers ITR ratios. Experience has shown that ITR ratios are fairly resistant to changes in the measurement environment. As a result, combining the ITR ratios observed for individual release transitions (as explained above) provides a reasonably good estimate of the ITR ratio that would result for a migration that spans all those releases.

The ITR ratios shown in Table 2 on page 22 are based on the following pairs of measurements:

Pair Number	From Runid	To Runid	Processor	Memory	CPU Util	Base Pg/cmd	ITR Ratio	Symbol
VM/SP 5 to VM/SP 6: FS7B0 Workload; Reference 1								
1	EC7620	EC4295	4381-13	16M	70	11	0.834	
2	EC7620	EC4295	4381-13	16M	80	15	0.811	
avg							0.82	(R5)
VM/SP 6 to VM/ESA 1.0 (370 Feature): FS7B0 Workload; Reference 1								
3	EC4295	EC7603	4381-13	16M	70	15	1.069	
4	EC4295	EC7603	4381-13	16M	80	20	1.075	
avg							1.07	(R6)
VM/XA 2.0 to VM/XA 2.1: FS7B0R Workload; Reference 1								
5	Y62R5401	Y6$R5401	3090-600J	512M/2G	90	15	1.02	(R20)
VM/XA 2.1 to VM/ESA 1.0 (ESA Feature): FS7B0R Workload; Reference 1								
6	Y2$R2001	Y23R2001	3090-200J	256M/2G	90	11	1.064	
7	Y6$R5401	Y63R5405	3090-600J	512M/2G	90	12	1.029	
avg							1.04	(R21)
VM/SP HPO 5 to VM/ESA 1.0 (ESA Feature): FS7B0R Workload; Reference 1								
8	Y25R1141	Y23R1143	3090-200J	64M/512M	90	22	0.97	(RH)
VM/ESA 1.0 (ESA Feature) to VM/ESA 1.1: FS7B0R Workload; Reference 2								
9	Y63R5866	Y64R5865	9021-720	512M/2G	90	13	1.059	
10	L23R1770	L24R1770	9121-480	192M/64M	90	14	1.032	
11	L13R0911	L14R0910	9121-320	192M/64M	90	12	1.045	
12	H13R0280	H14R0287	9221-170	48M/16M	80	11	1.043	
avg							1.04	(R1E)
VM/ESA 1.0 (370 Feature) to VM/ESA 1.0 (ESA): FS7B0R Workload; Reference 2								
13	H17R0281	H13R0280	9221-170	48M/16M	80	7	0.91	(R13B)
VM/ESA 1.0 (370 Feature) to VM/ESA 1.1: FS7B0R Workload; Reference 2								
14	H17R0281	H14R0287	9221-170	48M/16M	80	7	0.95	(R13C)

Table 3. VM Measurement Pairs

Explanation of columns:

Memory - The amount of real storage and (when applicable) expanded storage in the measured configuration.

CPU Util - Approximate processor utilization. The number of users is adjusted such that the “from” case runs at/near the stated utilization. The “to” case is then run at that same number of users.

Base Pg/cmd - The average number of paging operations per command measured for the base (“from”) case. This value gives an indication of how real memory constrained the environment is. For configurations with expanded storage used for paging, this value includes expanded storage in/out operations in addition to DASD page I/Os.

Symbol - Shows the symbol used to represent this release transition in Table 2 on page 22.

The results in this table illustrate the fact that the release-to-release ITR ratios can and do vary to some extent from one measured environment to another. For example, measurement pair 1 shows a somewhat better ITR ratio than measurement pair 2. This is because the first environment is less real storage constrained so the increase in real storage requirements that occurs in VM/SP 6 has less influence in that case.

A complete description of the measurement configuration and results for each of the VM measurement pairs shown above can be found in the following reports:

1. *VM/ESA Release 1.0 Performance Report*, ZZ05-0469[2]
2. *VM/ESA Release 1.1 Performance Report* (this document)

Table 3 on page 25 references this list to show where each measurement pair is documented.

Minidisk to Shared File System

Measurements were obtained on VM/ESA 1.1 to compare the performance of the CMS minidisk file system (EDF) to the Shared File System (SFS) to demonstrate the effects of migrating files from minidisk to SFS. Also, some VM/ESA 1.0 measurements were obtained as a reference for the previous release minidisk to SFS comparison. For these measurements all end user data was moved from minidisks to SFS. For a more complete description of these workloads, see "CMS Intensive (FS7B)" on page 273.

SFS requires more system resources (real storage, virtual I/O, processor busy time per command) than minidisk when performing similar tasks. The increase in processor busy time per command is slightly smaller in VM/ESA 1.1 than it was in VM/ESA 1.0.

SFS response times were similar to or somewhat longer than minidisk response times at similar processor utilization. With an equal number of users, processor utilization was significantly higher in the SFS case, resulting in much longer response times.

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

Note: A Coordinated Resource Recovery (CRR) server did exist for the SFS 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 SFS performance you must have a CRR server running, even in regression environments where you are not using CRR.

[2] This document is classified as IBM Internal Use Only. Contact you IBM representative for access to the information contained in this publication.

Virtual Machine Storage Considerations

Virtual machine storage is a consideration in migration and administration of VM/ESA. This includes both the storage size of virtual machines and the placement of saved segments. This section describes the impact to system resources required for different storage sizes of virtual machines and the placement of saved segments. The main issue is the real storage required by CP for the control blocks used in the management of virtual storage.

The two key types of control blocks are the page and segment tables. A segment table is created to represent the virtual machine storage. For a virtual machine with a storage size of 32M or smaller, the segment table fits inside the Virtual Machine Definition Block (VMDBK) which is 4K bytes in size. In the past, a limit of 31M existed when running VM/XA on a processor without the Storage Key Facility. For larger virtual machines, CP allocates a segment table with a size of 4K bytes per gigabyte of virtual machine storage. This larger segment table is located outside of the VMDBK and must be contiguous. In VM/XA, 999M is the maximum size of a virtual machine; VM/ESA 1.0 increased the limit to 2047M.

Page tables are kept in page management blocks (PGMBKs). Each PGMBK is 4K bytes in size and represents one megabyte of virtual machine storage. A PGMBK is created by CP when the corresponding megabyte of virtual machine storage is first referenced. The term "referenced" also includes storage key operations. Therefore, if the megabyte of storage is never referenced, CP never uses storage for a PGMBK describing it. PGMBKs are not created for megabyte gaps between DCSSs, or between the top of the virtual machine and the first DCSS. Starting in VM/ESA 1.0, PGMBKs are eligible to be paged out of main storage.

Prior to VM/ESA 1.1, PGMBKs were created for all of a virtual machine′s storage at IPL of CMS. This was due to storage key processing. As mentioned earlier, these PGMBKs were eligible to be paged out of real storage in VM/ESA 1.0. With the addition of the Pending Page Release function in VM/ESA 1.1, the storage key processing no longer requires the creation of PGMBKs for all of the virtual machine storage at IPL time. See "Pending Page Release" on page 7 for additional information.

The use of saved segments that are shared allows for a common set of PGMBKs to be shared among virtual machines. The segment tables for each virtual machine point to the same (shared) PGMBKs.

The considerations for the virtual machine storage size in VM/ESA are as follows:

1. They should be less than 32M to avoid the requirement for a separate segment table outside the VMDBK (additional 4K bytes).
2. Some applications or programs utilize virtual storage based on how much is available. This is an attempt to trade off greater storage requirements for less I/O processing.

The considerations for shared saved segment placement in VM/ESA are as follows:

1. Should be located below the 32M line to avoid the requirement for a separate segment table outside the VMDBK (additional 4K bytes). Only the PGMBKs are shared among virtual machines.
2. Some products and applications can not function above the 16M line.
3. Some products and applications do support 31 bit addressing and can be placed above the 16M line. However, some of these interact heavily with products that are not 31 bit addressable. In these scenarios, there can be additional costs associated with linkage between the two products. This cost is mostly in terms of CPU consumption.

Measurements were made to verify some of the items discussed above (see "Virtual Machine Storage Considerations" on page 111). A set of measurements dealt with the location of the CSL saved segment. Runs were made with CSL at 7M, 30M, and 35M. The 35M run did cause additional real storage requirements, as expected. However, paging was not constrained so the system configuration was able to absorb the increase without affecting external response time or ITR.

Software Mode Comparisons

The virtual mode comparison measurements were made to examine the performance effects of changing the user virtual machine modes. In VM/ESA 1.0, when going from 370 mode to XA mode, the processor busy time per command increased by 3.0% for the measured environment. In VM/ESA 1.1 this increase was reduced to 1.2%. This was accomplished by reducing the number of instructions executed in XA mode unique paths, primarily in the SVC interrupt handler. When running XC mode in VM/ESA 1.1, there was an additional 0.9% of CPU time required to support this new environment.

OfficeVision Migration from VM/XA 2.1

This section documents the migration data collected for an OV/VM environment. The base starting point was 6000 users running on VM/XA 2.1. Adequate performance was achieved at about 87% CPU utilization and an external response time of 0.980 seconds. A measurement was made increasing the users to 6200. This resulted in a very large increase in external response time and a reduction in the internal throughput rate, indicating that the system had become over loaded.

Using the 6000 user VM/XA 2.1 measurement as the base, VM was upgraded to VM/ESA 1.0. This environment had a positive effect on external response time, improving by 0.170 seconds (17%) with a slight decrease in the internal throughput rate. Another measurement was made increasing the number of users to 6200. This time the external response time only increased to 1.05 seconds (30%) and the internal throughput rate remained about the same, indicating that VM/ESA 1.0 could support this increased user load.

Using the 6200 user VM/ESA 1.0 measurement as the base, VM was upgraded to VM/ESA 1.1. An additional positive effect was observed on external response time, improving by 0.126 seconds (12%) with an ITR improvement of about 4.1%.

Both VM/ESA 1.0 and VM/ESA 1.1 have improved the performance of this OfficeVision environment. This is illustrated in the following figure.

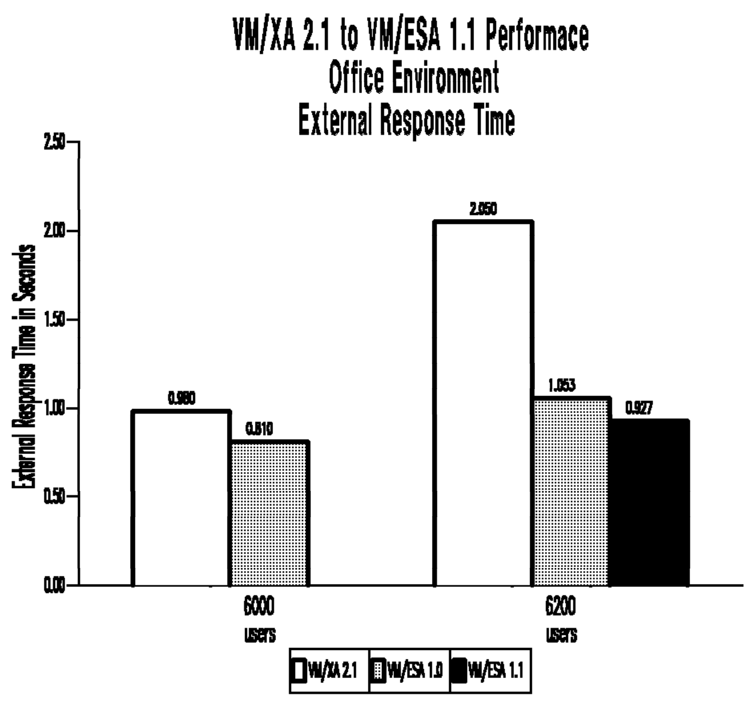

Figure 3. OfficeVision Migration from VM/XA 2.1

MVS Guest Migration from VM/ESA 1.0

Guest operating system performance on VM/ESA 1.1 was predicted to be equivalent to VM/ESA 1.0 because no functional changes were made to VM guest services. V=R MVS guest measurements verified that, in this instance, no performance changes had occurred for VM/ESA 1.1.

Figure 4 shows MVS workload (CB84) ITR as a percent of native ITR for the native, VM/ESA 1.0, and VM/ESA 1.1 guest environments.

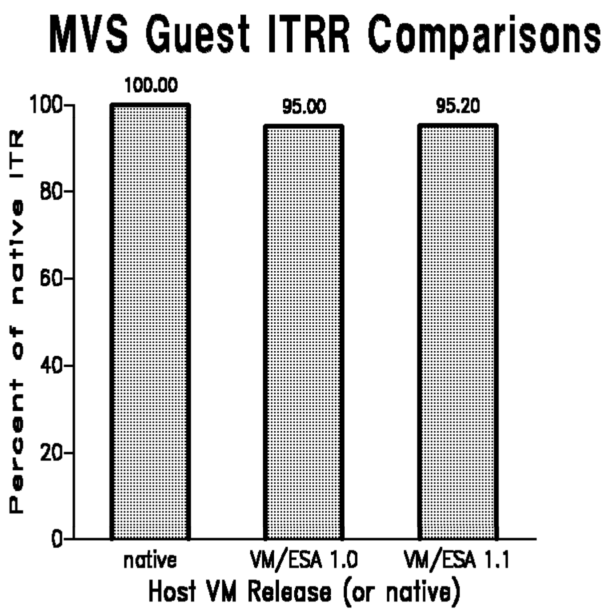

Figure 4. MVS Guest ITRR Comparisons

VSE Guest Migration from VM/ESA 1.0

Measurements indicated that, without APAR VM51012, VM/ESA 1.1 performance for VSE guests was equivalent to VM/ESA 1.0 ESA Feature and VM/XA 2.1. This is as expected since no major changes have been made to VM that would have a significant effect for VSE guests.

When APAR VM51012 was added to VM/ESA 1.1, however, the ITR improved dramatically for V=V VSE guests (i.e., those requiring DASD channel program translation by CP). The most significant improvement was for DASD I/O to minidisks although the improvement for dedicated DASD was dramatic as well. See "CCW Fast Path" on page 9 for an explanation of CCW fast path.

Other measurements quantified the behavior of a VSE batch system running as a guest of VM/ESA 1.1 in a dedicated LPAR. As expected, the V=R case degraded compared to the non-LPAR environment due to the lack of I/O assist and the need to perform CCW translation.

Figure 5 shows the VSE V=V MODE=ESA guest internal throughput rate (transactions per CPU busy minute) before and after CCW fast path APAR VM51012 is applied to the VM/VSE 1.1 host system.

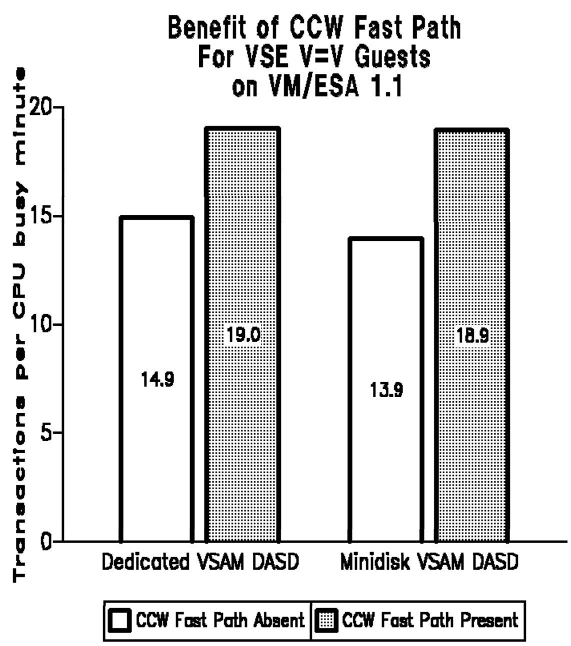

Figure 5. CCW Fast Path Benefit for VSE V=V Guest

3. Hardware Capacity

Processor Capacity

The processor capacity measurements were made to determine the performance of VM/ESA 1.1 when running on different size processors within the same family. The performance of VM/ESA 1.1 scaled as expected on these processors. For the 9021, expectations were based on similar measurements made on the 3090-600J using VM/ESA 1.0. For the 9121, expectations were based on the PD3 and HT4 hardware capacity workloads.

The following graph represents the internal throughput rate (ITR) as a function of the number of processors varied online for the 9021-720 and 9121-480.

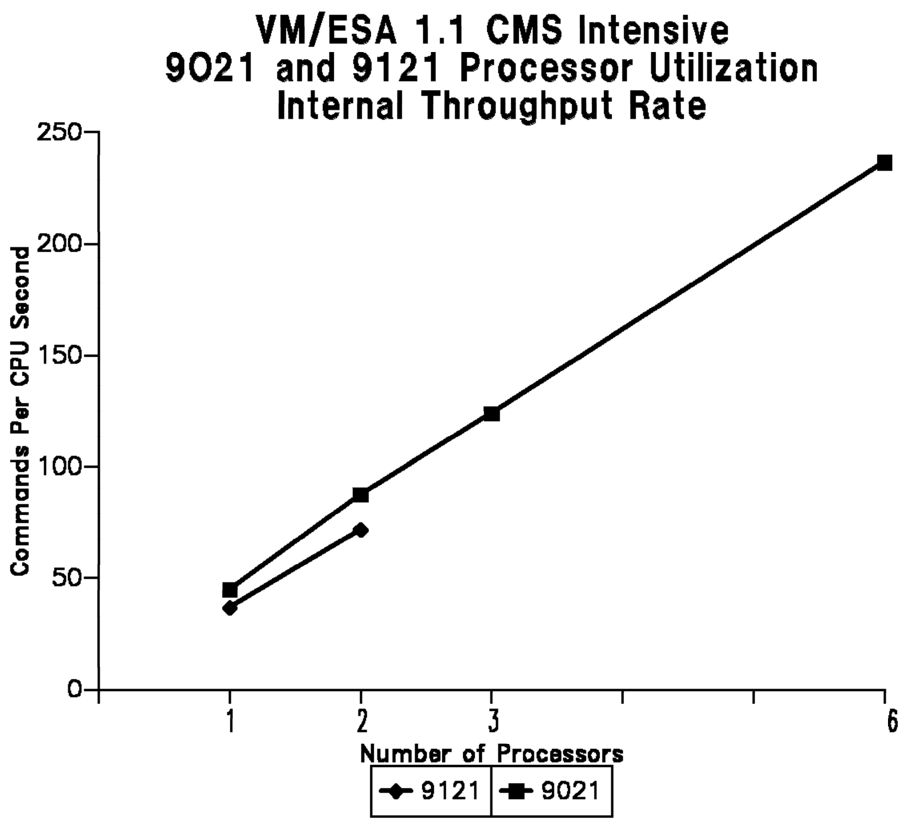

Figure 6. Internal Throughput Rate for Selected 9021 and 9121 Processors

Storage Constrained Runs

A set of six runs were completed on the 9021-720 processor to determine how well VM/ESA 1.1 performed in storage constrained environments and what the minimum storage requirements would be to run the FS7B35R workload with 4800 users and still achieve acceptable performance. All runs were completed with the same hardware and software configuration except for the real and expanded storage sizes.

The following two graphs show the measured external response times and paging rate per command for the various runs. The storage sizes used for each run are labeled along the X-axis of the graph, except for the minimum storage configuration that provided acceptable performance which is indicated by the dotted line. For the purposes of this discussion, performance will be considered "acceptable" when the external response time is less than one second. This is indicated in the response time graph by a horizontal dashed line at one second.

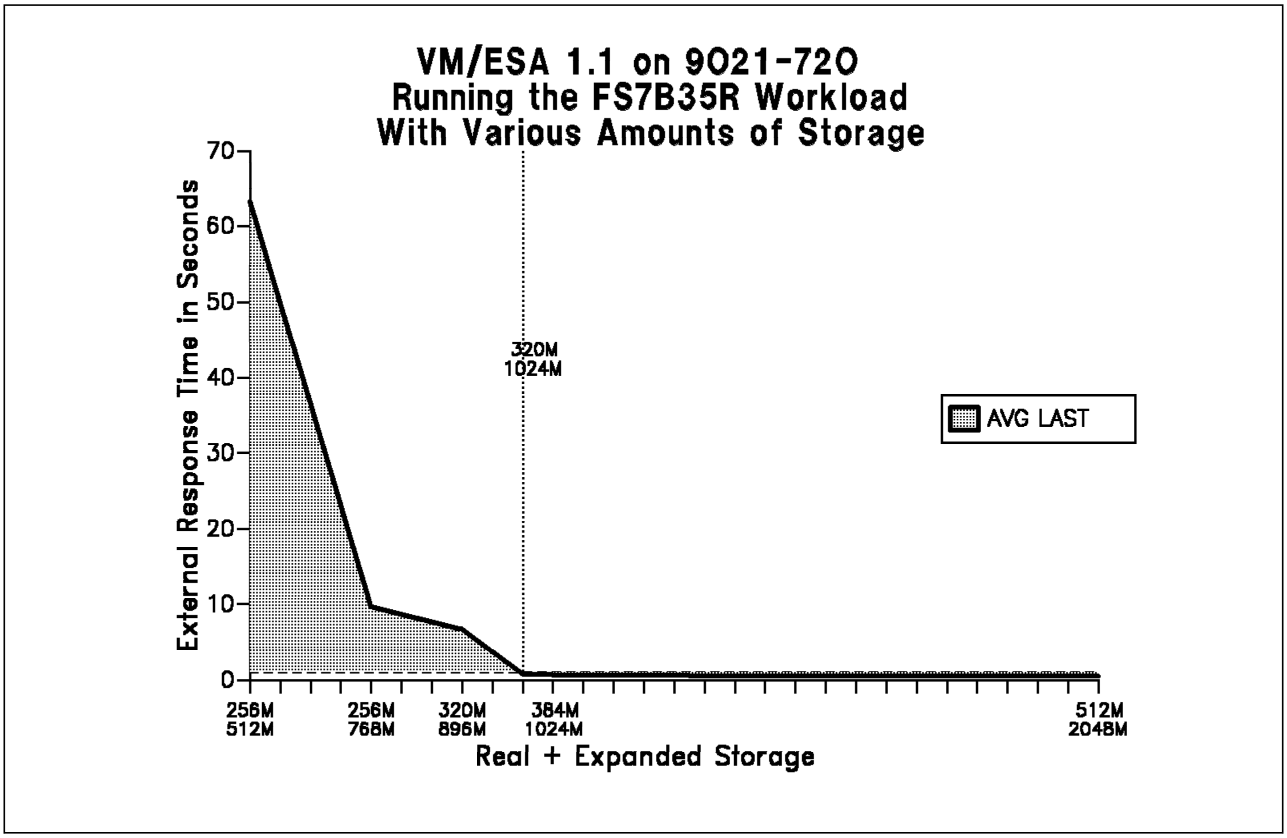

Figure 7. External Response Time for Storage Constrained Runs

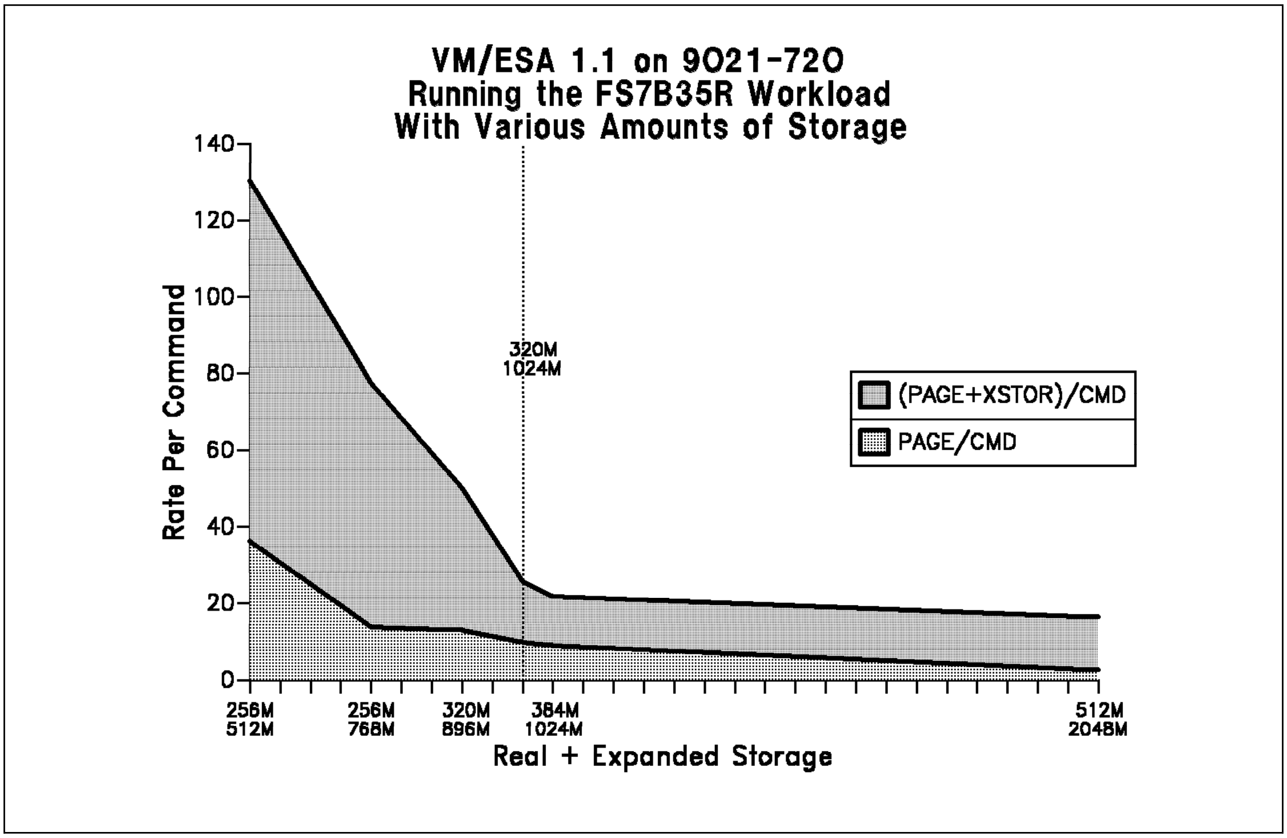

Figure 8. Paging Rates Per Command for Storage Constrained Runs

The graphs show that with decreasing storage size, external response times and paging rates increase sharply, and the 256M/512M run is clearly thrashing. Also depicted in the graph is the cutoff for acceptable response times for the FS7B35R workload, which was 320M/1024M. Taking away just 128M of expanded storage (320M/896M) caused external response times to jump over 6 seconds. Note that this minimum is applicable only to the FS7B35R workload; other workloads may require more or less storage.

The paging graph is split along the X-axis at the minimum storage size needed to provide acceptable response times. For all runs, the sum of the working sets of the logged on users does not fit in real storage so some paging always occurs. However, to the left of the line, when a virtual machine is running, some pages belonging to its working set are stolen. This forces the virtual machine to wait for page fault resolution. Proceeding from the split line to the Y-axis, more and more page stealing is occurring. To the right of the line, when a virtual machine is running, its pages are maintained in storage and not taken by the system. Of course, a virtual machine may still take page faults in this case, but not because the system is stealing pages from it. Proceeding from the split line to the right of the graph, less and less paging occurs. The more storage that is available to the system, the more likely the case that, when a virtual machine is re-dispatched, its pages are still in storage and it can continue running without any page fault delays.

4. New Functional Enhancements

VM Data Spaces

Exploitation by Shared File System

SFS exploits VM Data Spaces through read only access to DIRCONTROL directories. The data space contains a single shared copy of the File Status Tables (FSTs) and data for the files in the directory. Direct reference to the FSTs and file data is made when users are running in XC mode. When users are running in 370 or XA mode, a private copy of the FSTs is used and the file data is referenced indirectly via CP.

Comparing this environment to best case minidisk usage (minidisk caching and all read-only minidisks have their FSTs in shared segments) showed some increase in CPU usage and similar external response times. When compared to a typical minidisk environment (minidisk caching and one of four read-only minidisks has shared FSTs), the CPU usage and external response times are about the same.

When comparing SFS in data spaces to the typical minidisk environment without minidisk caching or cache controllers, the SFS case showed a decrease in external response time because the file blocks are cached in memory for the data space case.

When data spaces are exploited by 370 mode users, the FSTs are not shared. This resulted in a decrease in CPU usage but an increase in real storage requirements when compared to XC mode usage.

When data spaces are exploited, there was a significant decrease in CPU usage and paging compared to normal SFS usage.

Exploitation by Program Products

The SQL/DS 3.3 VM Data Spaces Support feature (VMDSS) makes use of VM data spaces in order to improve performance. Through the use of data spaces, dramatic reductions in response time and processor usage have been observed. Although VMDSS improves the performance of many types of SQL/DS requests, some of the largest improvements have been for query requests in environments that are not memory constrained. Early VMDSS performance results can be found in *SQL/DS VMDSS Presentation Guide*.[3]

VS FORTRAN Version 2 Release 5.0 allows users to address larger data areas via a new function, called "Extended Common," by placing these blocks in VM Data Spaces. Extended Common allows each dynamically allocated common block to be as large as 2GB. Support for VM/ESA 1.1 became available in December 1991.

3 See the Related Publications section at the end of this book for information on how to obtain this publication.

3990-3 DASD Fast Write

DASD fast write is a 3990 feature which decreases write response time by immediately returning channel end and device end when a write hit occurs. Then, the 3990 controller does the actual write asynchronously when the device is available. The data is saved in nonvolatile storage (NVS), thereby eliminating the possibility of data loss even if a power failure occurs.

The performance benefits of DASD fast write were evaluated in both CMS intensive and OfficeVision environments. The measurements were made using DASD configurations with various read:write ratios. Both environments showed large improvements in DASD response times and DASD utilization which, in turn, resulted in improved total system responsiveness.

The CMS intensive results (FS7B0R workload) showed the following:

- For user minidisk volumes, average DASD response time decreased by 51% to 65%, while average device utilization decreased by about 45%.
- For the spool volumes, average DASD response time decreased by 30%, while average device utilization decreased by 22%.
- Total system external response time improved by 6%. Internal response time decreased by 11% to 14%.
- The number of DASD actuators can be reduced while preserving, or even improving, system responsiveness. For the cases examined, going from 14 minidisk volumes without DASD fast write to 6 minidisk volumes with DASD fast write resulted in a net 3.6% decrease in external response time.
- It is best to exclude volumes that experience very high rates of minidisk disk formatting from DASD fast write eligibility.

The OfficeVision results (IOB workload) showed that:

- For the Calendar machine minidisk volume, average DASD response time decreased by 43%, while average device utilization decreased by about 44%. This should allow for an increase in capacity of between 1.4 to 1.6 times as much calendar activity.
- For the Database machine minidisk volume, average DASD response time decreased by 64%, while average device utilization decreased by about 65%. This should allow for an increase in capacity of between 1.8 to 2.2 times as much database activity.
- For the Mailbox machine minidisk volumes, average DASD response time decreased by 77%, while average device utilization decreased by about 72%. This should reduce the number of Mailbox machines required.
- For all DASD fast write volumes used, average DASD response time decreased by 75%, while average device utilization decreased by about 68%.
- Total system external response time improved by 23%. Internal response time decreased by 25%.

In conclusion, DASD fast write can be very effective at improving I/O subsystem performance. This can be used to benefit the overall system in a number of ways:

- Improving system responsiveness.

- Increasing system capacity.
- Removing or avoiding server bottlenecks.
- Reducing requirements for multiple servers.
- Reducing the number of DASD actuators that are required.

For measurement results and further discussion, see “9021-580 / Minidisk” on page 180 and “9021-580 / OfficeVision” on page 188. See “Related Publications” on page 343 for a list of related publications.

CMS Pipelines

CMS Pipelines is now included in the VM/ESA 1.1 product. CMS Pipelines allows a user to direct the output of one program as input to another with the CMS command PIPE. In addition to connecting programs together, CMS Pipelines provides users with a number of built-in filters to manipulate the data being passed between the programs and allows users to write their own filters. CMS Pipelines also has a REXX interface so pipeline statements can be put in REXX EXECs and the results of a pipeline statement can be stored in REXX variables.

Comparison to PRPQ 1.1.6 CMS Pipelines

Performance tests were made to ensure that the performance of VM/ESA 1.1 CMS Pipelines was equivalent to the currently available PRPQ version 1.1.6. A half dozen test cases were developed that included many of the different features of CMS Pipelines including filters, device drivers, issuing CP and CMS commands, and multi-stream pipelines.

In terms of the key performance indicators of estimated CPU time, pathlength, storage use and privileged operations use, VM/ESA 1.1 CMS Pipelines was found to be equivalent to PRPQ 1.1.6 CMS Pipelines. In fact, VM/ESA 1.1. CMS Pipelines was slightly better on average (less than 1%) than PRPQ 1.1.6 in terms of estimated virtual CPU time per test case.

REXX/EXECIO/XEDIT vs. CMS Pipelines

In addition to the comparison to PRPQ 1.1.6, VM/ESA 1.1 CMS Pipelines was compared to REXX/EXECIO/XEDIT. Twelve “functions” were coded in REXX/EXECIO/XEDIT and with CMS Pipelines. These functions were various I/O and data manipulation tasks that were coded in previously available methods (such as REXX loops, EXECIO and XEDIT macros) and with pipeline statements. See “CMS Pipelines” on page 196 and “CMS Pipelines” on page 290 for a more complete description of these test cases.

It was found that CMS Pipelines performed better than its REXX/EXECIO/XEDIT equivalents when filters were used to manipulate the data rather than using REXX loops and program statements and/or XEDIT macros. Estimated virtual CPU times for the CMS Pipelines versions of these types of tests were in some cases 80% less than the REXX/EXECIO/XEDIT versions. Not only did it perform better, CMS Pipelines also saves in lines of code needed to implement these (and similar) functions.

It was found that CMS Pipelines performed worse than its REXX/EXECIO/XEDIT equivalents for the test cases where CMS Pipelines filters were not used or needed. An example of this may be the reading of the entire contents of a small file into storage. This can be accomplished with a single EXECIO statement or CMS Pipelines statement. There is not much of an advantage of using CMS Pipelines over EXECIO in this case and the performance (in terms of CPU time and pathlength) was found to be worse.

For all test cases, CMS Pipelines used less non-shared storage per user than the REXX/EXECIO/XEDIT equivalents when installed in a shared segment. CMS Pipelines also issued fewer unassisted privileged operations per test case on average than did REXX/EXECIO/XEDIT.

In conclusion, when the input/output data of a program is to be manipulated, and pipeline filters can be used to do this rather than REXX/EXECIO loops or XEDIT macros, the performance of the CMS Pipelines implementation is likely to be better, and in some cases significantly better.

GCS IPOLL Option

In VM/ESA 1.1, GCS offers a new tuning option called IPOLL. It can be used to improve IUCV request handling efficiency for GCS applications, such as VSCS, that support IPOLL ON. See “GCS IPOLL” on page 10 for further discussion of the IPOLL option.

Measurements were obtained that assess the performance effects of this option in the CMS intensive and OfficeVision environments. Mixed results were obtained. For both environments, use of IPOLL ON did result in the expected decrease in IUCV requests and that, in turn, resulted in a decrease in CP CPU usage (-0.55% for CMS intensive and -0.36% for OfficeVision). However, in the CMS intensive case, there was an overall slight decrease in total CPU usage and no discernible effect on response time while, in the OfficeVision case, there was an overall 1% increase in CPU usage and a 0.096 second (10%) reduction in external response time.

In conclusion, although IPOLL ON tends to have a positive effect on performance, its effects are typically small and tend to vary from one environment to another.

Inter-System Facility for Communication (ISFC)

ISFC is a new function in VM/ESA 1.1 systems that provides high-speed connectivity to groups of LAN attached workstations running VM PWSCS. The ISFC function is implemented directly in the Control Program (CP). This design allows for greater communications throughput by eliminating the communications server virtual machine. By eliminating the server virtual machine, which acted as a sort of "middleman" to connect Virtual Machine communications partners, overhead is reduced and throughput is increased.

Another feature provided by ISFC is collection management. ISFC allows for the dynamic formation of Communication Services (CS) collections. These collections may consist of either a VM/ESA 1.1 system or a VM/SP 6 system running the VM PWSCF PRPQ, and LAN-attached workstations running any of the following operating system environments: OS/2, Windows, DOS, Novell Netware, AIX 1.2 and AIX 3.1. This report documents the performance of a communications triad composed of a VM/ESA 1.1 system, an OS/2 system configured as a domain controller, and an OS/2 system configured as a user workstation.

The following observations were made regarding the performance of CS collections:

- Throughput was usually better when data was sent and received in large chunks. This practice helped decrease the number of API crossings through the protocol stack.
- Analysis of resource consumption on the host (a 3090-300J processor) showed that only four to five percent of the CPU was utilized. There was potential for greater throughput; the host was definitely not a source for performance bottlenecks.
- Performance increased when communications adapters and software were configured as follows:
  - The System/370 Channel Adapter /A was configured above the 1M line. This change was made using the backup copy of the reference diskette. The backup copy of the reference diskette must have the code from the option diskette shipped with the adapter copied onto it.
  - The token ring adapter was set to run at 16 Mbps., the fastest setting allowed. In order for the token ring to function properly, each system on the LAN was set to run at the faster data rate.
  - The Communications Manager advanced configuration was used to configure the largest permitted transmit buffer size. This varies with the data rate. When the token ring was running at 4 Mbps., the transmit buffer size was set to 4 Kb. When the token ring was running at 16 Mbps. the transmit buffer size was set to 8 Kb.

ECKD-Formatted DASD versus CKD-Formatted DASD

The objective of these measurements is to ensure similar performance across 3380 system paging devices when these devices are in ECKD-format versus the CKD-format. Even when run at a high DASD utilization, the paging rates and page DASD response times for the measured system showed no significant performance effect going from CKD-formatted paging devices to ECKD-formatted paging devices.

5. Tuning Considerations

Recommended 9221 Tuning

The tuning recommendations, specifically for the 9221, concentrate on reducing I/O instructions and SIE instructions. These tuning recommendations significantly improved internal throughput and response time. Listed below are the tuning options. For more details, see section "Recommended 9221 Tuning" on page 223.

- Minidisk Cache

  Configure a portion of real storage as expanded storage and use it exclusively for minidisk cache. The result is a replacement of DASD I/Os with less CPU intensive minidisk cache reads.

- DSPSLICE

  Increase the default dispatch slice to three times the default. This reduces in the number of timer interrupts for time slice end processing and the associated SIE instructions.

- VTAM Delay

  Set the VTAM delay to 0.2. This reduces VTAM I/O and the associated SIE instructions.

- IPOLL ON

  Set IPOLL ON for VTAM. This reduces the number of IUCV instructions and the associated SIE instructions.

- Preloaded Shared Segments

  Load the FORTRAN and Script shared segments from an idle user during system startup. This prevents the shared segment′s page frames from becoming invalid when not in use and avoids page reads when the next user wants to access the shared segments.

Using XSTOR on a 9121

On 9121 processors, the installation may take some of the real storage and use it as expanded storage when no true expanded storage is installed on the machine. This raises the question of whether it is better to take some of real storage and use it as expanded storage or run without any expanded storage at all. The assumption was made that it would not be advantageous to use real storage as expanded storage used for paging. If the storage can be used as real storage, it should be more beneficial to use it that way so as to decrease paging. However, it was unclear whether it would be more beneficial to use some expanded storage exclusively for minidisk caching, or to have no expanded storage at all and forego minidisk caching.

Two runs were completed on a 9121-480 with 256M real storage. The first run had no expanded storage at all while the second run had 64M (the minimum amount that can be taken) of real storage used as expanded storage exclusively for minidisk caching. The results of these runs indicated that for the FS7B35R workload, using a portion of real storage as expanded storage for minidisk caching performed slightly better than using no expanded storage at all.

It should be noted here that 256M of real storage with no true expanded storage installed on the machine was a little tight when trying to run the FS7B35R workload at a processor utilization of 90%. If true expanded storage had been available, using a portion for minidisk caching would have been an obvious choice to help reduce I/O times. If storage is constrained, (such as it was with the FS7B35R workload on the 9121-480), installations may not gain in performance as the FS7B35R workload did. FS7B is a very uniform workload that experiences an especially high minidisk cache hit ratio. Therefore, it gains disproportionately from replacing more expensive minidisk I/Os with faster page I/Os. Other more storage intensive and less I/O intensive or less uniform workloads may benefit more by using no minidisk caching and having more storage for paging.

Set Reserve Option

Initial FS7B35R runs made with VM/ESA 1.0 on the 9121-480 with 256M real storage indicated that performance was unacceptable (external response times as large as 36 seconds). A close examination of the performance data indicated that serial page faulting was occurring in the SFS and VSCS servers. What was happening was that the SFS and VSCS servers were often in page wait, in effect serializing the servers and causing all the dependent users′ response times to degrade.

It was decided to reserve the servers′ working sets in storage with the SET RESERVE command. This tuning option dramatically improved performance. External response time was reduced from 36 seconds to less than one second. The amount of time the servers spent in page wait was greatly reduced as was the system paging rate. See “Set Reserve Option” on page 232 for more details on the performance of these runs.

There is a potential down side to using the SET RESERVE command. Reserving pages for a given user may cause other users to experience increased paging due to fewer pages left in the Dynamic Paging Area (DPA). Care must be taken not to reserve more pages than are needed by the virtual machine. SET RESERVE should most often be used for those virtual machines that, when taking a page fault, will degrade the performance of more than just that particular virtual machine. Examples include servers and guest operating systems.

OfficeVision MSGFLAGS Settings

This section documents the tuning data collected for an OV/VM environment. Using the OV/VM MSGFLAGS command for both the Mailbox and the Calendar server machine, a set of measurements were made to see what effect these messages have on system capacity in this 6-way environment. First, a comparison between messages on (default) and messages off, at the same number of users was performed. When messages were turned off, there was a reduction of 0.06 seconds (7%) in external response time and a 3% reduction in CPU busy time. Further, an increase in the percent of emulation time on the master processor implied that more user work was allowed to run on the master processor.

Turning messages off allowed for improved capacity due to reduced resources consumed and potentially reduced master processor requirements.

Part 3. Specific Measurements

6. Introduction

This part of the report contains the configuration details and specific results for those measurements obtained to evaluate the performance of VM/ESA 1.1.

Format Description

For each group of measurements there are five sections:

1. WORKLOAD: This section specifies the name of the workload associated with the data. For more detail see Appendix C, “Workloads” on page 273.
2. HARDWARE CONFIGURATION: This section summarizes 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 used on the processor.
   - DASD: The DASD configuration used during the measurement.

     The SYSTEM volumes are those where the operating system code and data areas reside.

     The remaining volumes detail the DASD used to run the measurement. The table indicates the type of volumes used, the number and type of control units which connect these volumes to the system, and the distribution of these volumes. The PAGE, SPOOL, TDISK, USER, and SERVER headings indicate how many full volumes were used for system paging, spooling, temporary disk space, user minidisks, and server minidisks respectively.
   - TAPE: The tapes being used and what they were used for.
   - COMMUNICATIONS: The type of controller, the number of communication controllers, the number of lines per controller used for the measurement, and the line speed. Since the 3745-410 controllers used for this report are split in half and run in twin-dual mode, the number specified is the number of halves used for each of the processors (i.e. the driver and the processor being measured). Each of these halves has a maximum of 50 lines available and can support a maximum of 3000 users.
3. SOFTWARE CONFIGURATION: This section contains pertinent software information. Based on the type of measurement, this section contains a subset of the following subsections:
   - DRIVER: The tool used to simulate users.
   - THINK TIME DISTR: The type of distribution used for the user think times.

     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. This type of mechanism allows the transaction rate to vary depending on the command response times in the measurement.

IOB think time distribution represents the think time dictated by the IBM Office Benchmark (IOB V2.1) workload. The think time includes an average two second delay between commands issued by TPNS, the built in think times which are part of the IOB scripts, and the IOB script scheduling algorithm. When users finish executing a script, the script scheduling algorithm calculates how much time was spent executing the script, subtracts this number from ten minutes, and delays the user for the resulting amount of time. Thus if a script was executed in 7.9 minutes, the user would be delayed for 2.1 minutes before starting the next script and this time would be included in the user's think time. This has a tendency to keep the message rate per user constant across all of the measurements.

- CMS BLOCKSIZE: The blocksize of CMS minidisks.
- USER VM SIZE: The storage size of the user virtual machine.
- USER CMS MODE: The software machine mode (370, XA or XC) 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: The name and type of the server machines, their storage size, their software machine mode, their scheduling relative share, and any special options used.
- MVS VERSION: The software version of MVS used for the MVS measurements.
- V=R SIZE: The size of the V=R area for the V=R guest measurements.
- GUEST VM SIZE: The storage size of the guest virtual machine.
- GUEST MACHINE MODE: The software machine mode (370, XA or ESA) of the guest virtual machine.
- OPERATING SYSTEM: The operating system on the processor for the connectivity measurements.
- SOFTWARE: The software being used for the connectivity measurements.
- SYSTEM MEMORY: The amount of memory available when the processor is a PC.
- FIXED DISK: The size of the fixed disk when the processor is a PC.
- CHANNEL ADAPTOR: The hardware card which enables a PS/2 processor to connect and communicate with the host via a 3088 controller.
- COMMUNICATIONS ADAPTOR: The hardware card which enables the PC to communicate with other PCs over a LAN.
- TRANSMIT BUFFER SIZE: The Communication Manager setting that controls the amount of data which may be sent over a LAN communications adapter in a single send. The Communications Manager is a component of OS/2 Extended Edition.

- PWSCS CONFIGURATION: The settings for the PWSCS tuning variables for PC processors using PWSCS.

4. MEASUREMENT DISCUSSION: This section contains additional information explaining the measurements that were done. It also contains an analysis of the performance data in the table and gives the overall performance findings.

5. Measurement Data: This section contains the table of performance results. This data was obtained or derived from the tools listed in the next section.

   There are several cases where the same information is reported from two sources because the sources calculate the value in a slightly different manner. As an example, consider the external throughput rate measures, ETR (T) and ETR, which are based on the command rate calculated by TPNS and RTM respectively. TPNS is external to the system and can directly count the command rate as it executes the commands in the scripts. Since CP is internal to the system, it has to make assumptions as to when transactions begin and end. This can make the counts reported by RTM vary in meaning from run to run and vary from the values reported by TPNS. As a result, the analysis of the data relies more on values using the TPNS command rate. Further, some values in the table (like TOT INT ADJ) have been normalized to the TPNS command rate in an effort to get the most accurate performance measures possible.

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

Tools Description

A variety of program products and tools were used in executing and evaluating the performance measurements. The program products which were used in conjunction with the measurements in this report (and are available to customers) are listed below:

RTM (Real Time Monitor) provides on-line performance analysis and determination facilities for VM systems. Since RTM is modified for each release of VM, this report contains data from VM/XA RTM/SF 1.4 for VM/XA systems, RTM VM/ESA 1.5 for VM/ESA 1.0 ESA systems, and RTM VM/ESA 1.5.1 for VM/ESA 1.1 systems.

VMPRF (VM Performance Reporting Facility) is the VM Monitor reduction program.

VMMAP (VM Monitor Analysis Program) is the VM/370 Montior reduction program.

VMPAF (VM Performance Analysis Facility) is a tool for performance analysis of VM systems.

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

RMF Monitors and reports MVS performance.

The internal tools which were used in conjunction with the measurements in this report are listed below:

Hardware Monitor Collects branch, event, and timing data.

VUMAPC Reduces hardware monitor data.

FSTTAPE Reduces hardware monitor data.

MONFAST Collects branch, event, and timing data on a 9221 in addition to reducing the data it collects.

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

REDFP Consolidates the QUERY FILEPOOL STATUS data.

Since each workload used a different subset of the tools and program products, the tools and programs used are itemized below:

- For the CMS intensive (FS7B) measurements: RTM, VMPRF, the hardware monitor, VUMAPC, TPNS, and TPNS Reduction Program. For measurements that included the Shared File System, the QUERY FILEPOOL STATUS command and REDFP were used as well. For the 9221 measurements, MONFAST was used instead of the hardware monitor and VUMAPC. For the VM/ESA 1.0 370 Feature 9221 measurements, VMMAP was used instead of VMPRF.
- For the OfficeVision (IOB) measurements: RTM, VMPRF, the hardware monitor, VUMAPC, VMPAF, TPNS, and TPNS Reduction Program,
- For the MVS Guest (CB84) measurements: RTM, the hardware monitor, and RMF.
- For the VSE Guest (PACEX8) measurements: RTM, VMPRF, the hardware monitor, and FSTTAPE.
- For the CMS Pipelines measurements: RTM, VMPRF, the hardware monitor, VUMAPC, TPNS, and TPNS Reduction Program.
- For the INSTVER Communications measurements: RTM and VMPRF.

7. Migration/Regression

CMS Intensive Migration from VM/ESA 1.0

For the following 9021 and 9121 regression measurements of the CMS intensive environment, the methodology used to determine the number of users required was to run VM/ESA 1.0 at 90% processor utilization. Then VM/ESA 1.1 was run with the same number of users.

For the 9021 processor runs, the RETAIN XSTORE MDC command was used to retain 64M of expanded storage as the minimum amount for minidisk caching. On the 9121 processors, which had no true expanded storage available, expanded storage was created by allocating real storage for this purpose. For the 9121 processor runs, all of the expanded storage was reserved for minidisk caching. See “Using XSTOR on a 9121” on page 43 for more information.

For the 9221 processor runs, the methodology used to determine the number of users required was to run VM/ESA 1.0 370 Feature at 80% processor utilization. VM/ESA 1.0 and VM/ESA 1.1 were run with the same number of users and with explicit tuning options. See section “Recommended 9221 Tuning” on page 223 for more information.

The 9221 showed an average think time of around 28 seconds versus approximately 26 seconds on the 9021 and 9121. On the 9221, the measurement period was one hour versus 30 minutes on the 9021 and 9121. The longer the measurement period, the closer the average think time approached the 30 seconds defined by the Bactrian distribution.

When comparing the VMMAP data shown for the 370 Feature measurements to the corresponding VMPRF and RTM data shown for the VM/ESA 1.0 ESA Feature and VM/ESA 1.1 measurements be aware that many of these measures do not have precisely the same meaning.

9021-720 / Minidisk

Following is a description of the environment used for the minidisk regression measurements on the 9021-720 comparing VM/ESA 1.0 and VM/ESA 1.1.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

- PROCESSOR: 9021-720
- STORAGE:
  - RSTOR: 512M
  - XSTOR: 2048M
- DASD:

	PACK NAME	TYPE
- SYSTEM:	PSYS02	3380-A
	PSPT01	3380-D
	WKLD01	3380-D
	WKLD02	3380-D

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
		NUMBER OF PACKS				
3380-A	15 - 3880-3	20	8	12	20	0
3380-D	3 - 3880-3	0	0	0	20	0

- TAPE: MONITOR 3480
- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	44	56Kb

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2M
- USER CMS MODE: XA
- USER RELSHARE: 100
- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON

4) MEASUREMENT DISCUSSION

The following table summarizes the results of the regression measurements on the 9021-720 comparing VM/ESA 1.0 and VM/ESA 1.1 for the minidisk-only CMS intensive environment.

Differences can be seen in the CPU resource utilization in this environment for the two releases. Overall processor requirements (PBT/CMD (H)) decreased by 5.5%, resulting in a 5.9% increase in internal throughput (ITR (H)). Almost all of the decrease in CPU time was in CP (CP/CMD (H)), which was down by approximately 13%. This drop in CP CPU time was due largely to the following five performance improvements:

- Pending Page Release

  The Pending Page Release changes also account for the decrease seen in the number of DIAGNOSE X¢10¢ instructions executed per command (DIAG 10/CMD) and the increase of DIAGNOSE X¢214¢ instructions (DIAG 214/CMD).

- CP Fast Dynamic Linkage
- IUCV Improvements
- XA Mode Improvements
- MDC Spin Lock Fix

  The MDC Spin Lock change accounted for at least 2 percentage points of the 5.8% increase in internal throughput. These changes remove a large system effect seen on high-end, n-way processors having large amounts of expanded storage used for minidisk caching, as in this environment.

See chapter 1, "Changes That Affect Performance" on page 7 for more detail on these and other performance improvements for VM/ESA 1.1.

Real storage requirements increased somewhat for VM/ESA 1.1 as evidenced by the overall increase in paging (PAGE/CMD plus XSTOR/CMD). Some of the growth in real storage requirements was in CP, as evidenced by the decrease in the number of pageable pages available (PGBLPGS) and the growth in the storage required for CP control blocks (FREEPGS). User working set size (WKSET (V)) grew by 7%, reflecting growth in CMS storage usage.

External response time (AVG LAST (T)) decreased by 0.290 seconds (35%) while internal response time (TOT INT ADJ) decreased by 0.027 seconds (11%). Any changes to the three basic system resource measures of CPU usage, real storage requirements, and I/Os can affect response time values. In this environment, the benefits of reduced CPU usage greatly outweighed the effects of the increased real storage requirements, resulting in the significant net improvement in response time.

The external response time improvement was much greater than that for the internal response time. It appears that the IUCV improvements, in addition to saving pathlength in general, also allowed VTAM to be more responsive in handling message traffic.

In summary, the performance of VM/ESA 1.1 showed much improvement over that of VM/ESA 1.0 for the minidisk-only CMS intensive environment. This environment showed lower response times and increased processor capacity.

RELEASE RUN ID	ESA 1.0 Y63R5866	ESA 1.1 Y64R5865
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**5860**	**5860**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
Response Time		
TRIV INT	0.057	0.056
NONTRIV INT	0.328	0.303
TOT INT	0.192	0.193
TOT INT ADJ	0.251	0.224
AVG FIRST (T)	0.633	0.397
AVG LAST (T)	0.817	0.527
Throughput		
AVG THINK (T)	26.00	25.19
ETR	266.53	238.56
ETR (T)	204.06	205.97
ETR RATIO	1.306	1.158
ITR (H)	223.28	236.35
ITR	48.61	45.71
EMUL ITR	78.49	69.86
ITRR (H)	1.000	1.059
ITRR	1.000	0.940
Proc. Usage		
PBT/CMD (H)	26.872	25.386
PBT/CMD	26.855	25.344
CP/CMD (H)	10.640	9.268
CP/CMD	10.193	8.739
EMUL/CMD (H)	16.229	16.116
EMUL/CMD	16.662	16.604
Processor Util.		
TOTAL (H)	548.36	522.88
TOTAL	548.00	522.00
UTIL/PROC (H)	91.39	87.15
UTIL/PROC	91.33	87.00
TVR(H)	1.66	1.58
TVR	1.61	1.53
Storage		
WKSET (V)	54	58
PGBLPGS	105K	104K
PGBLPGS/USER	18.3	18.2
FREEPGS	13897	14349
FREE UTIL	0.96	0.96
SHRPGS	1195	1186
Paging		
READS/SEC	266	332
WRITES/SEC	161	186
PAGE/CMD	2.093	2.515
XSTOR IN/SEC	1053	1053
XSTOR OUT/SEC	1267	1298
XSTOR/CMD	11.369	11.414
FAST CLR/CMD	8.189	5.651

RELEASE RUN ID	ESA 1.0 Y63R5866	ESA 1.1 Y64R5865
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**5860**	**5860**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
I/O		
VIO RATE	1702	1751
VIO/CMD	8.341	8.501
MDC READS	1183	1223
MDC WRITES	536	550
MDC MODS	452	463
MDC HIT RATIO	0.93	0.93
PRIVOPs		
PRIVOP/CMD	19.015	20.075
DIAG/CMD	17.098	23.603
DIAG 08/CMD	0.720	0.719
DIAG 10/CMD	5.709	0.015
DIAG 58/CMD	1.225	1.219
DIAG 98/CMD	0.309	0.291
DIAG A4/CMD	3.945	4.083
DIAG A8/CMD	1.872	1.893
DIAG 214/CMD	na	12.371
SIE/CMD	65.235	54.688
SIE INTCPT/CMD	42.403	35.547
FREE TOTL/CMD	100.362	84.517
VTAM Machines		
WKSET (V)	1955	1791
TOT CPU/CMD (V)	4.3505	3.8211
CP CPU/CMD (V)	2.4394	1.8823
VIRT CPU/CMD (V)	1.9112	1.9388
DIAG 98/CMD (V)	0.313	0.295
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 4. Minidisk Regression from VM/ESA 1.0 to VM/ESA 1.1 on the 9021-720.

9021-720 / 35% SFS

This section discusses the SFS regression measurements on the 9021-720 running VM/ESA 1.0 and VM/ESA 1.1.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-720
- STORAGE:
   - RSTOR:        512M
   - XSTOR:       2048M

- DASD:
            PACK NAME  TYPE
   - SYSTEM: PSYS02    3380-A
             PSPT01    3380-D
             WKLD01    3380-D
             WKLD02    3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-A	10 - 3880-3	20	8	12	0	0
3380-K	4 - 3990-2	0	0	0	0	16

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	22	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR:  BACTRIAN
- CMS BLOCKSIZE:     4K
- USER VM SIZE:      2M
- USER CMS MODE:     XA
- USER RELSHARE:     100
```

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
SERVE2	SFS	32M/XA	1500	QUICKDSP ON
SERVE4	SFS	32M/XA	1500	QUICKDSP ON
SERVE7	SFS	32M/XA	1500	QUICKDSP ON
SERVE8	SFS	32M/XA	1500	QUICKDSP ON
CRRSERVA	CRR	16M/XA	100	

4) MEASUREMENT DISCUSSION

The performance of VM/ESA 1.1 showed improvement over that of VM/ESA 1.0. The 35% SFS workload showed increased ITR, lower response times, and reduced processor utilization while maintaining the same command rate.

There was a 6.6% increase in ITR from VM/ESA 1.0 to VM/ESA 1.1. This was due mainly to a decrease in CP overhead as discussed in "9021-720 / Minidisk" on page 54. Of the CP performance improvements, the IUCV improvements had a greater impact on response time for the SFS workload than the minidisk workload since SFS uses APPC/VM to communicate with users and the *BLOCKIO interface for its I/O operations. This was evidenced by a greater drop in CP/CMD for the SFS case than appeared in the minidisk-only measurements. In the SFS regression measurements, CP/CMD decreased by 14.5% while the decrease for the minidisk-only measurements was 12.9%. This also had greater impact on response time for the SFS case. The decrease in CP/CMD was also reflected in lower processor utilization and a lower total:virtual ratio (TVR (H)) as well as reduced overhead in the VTAM and SFS service machines.

Both internal and external response time improved from VM/ESA 1.0 to VM/ESA 1.1. Internal response time (TOT INT ADJ) decreased by 0.071 seconds (22.5%) in VM/ESA 1.1. This can be attributed to:

- A decrease in CP overhead reflected in CP/CMD as mentioned above.
- CMS improvements mentioned in "9021-720 / Minidisk" on page 54.
- The SFS Checkpoint Improvement. The SFS Checkpoint Improvement accounts for most of the larger reduction in internal response time (TOT INT ADJ) for these measurements as compared to the minidisk-only measurements in the previous section. The SFS Checkpoint Improvements were responsible for a decrease of 0.039 seconds/command (-30.7%) in SFS TIME/CMD (Q). This improvement is discussed in more detail in chapter 1, "Changes That Affect Performance" on page 7.

External response time decreased by 0.237 (31.9%) seconds from VM/ESA 1.0 to VM/ESA 1.1. About 30% of this decrease in external response time (AVG LAST (T)) can be accounted for by the decrease in internal response time. The majority of the remaining response time reduction was due to a decrease in overhead in the VTAM machines. This was primarily due to the IUCV improvements that were incorporated into VM/ESA 1.1 as mentioned in the previous section.

Real storage requirements grew in VM/ESA 1.1 for reasons discussed in the previous section.

The increase in DIAG/CMD is due to the replacement of DIAGNOSE X¢10¢ instructions with DIAGNOSE X¢214¢ instructions for the pending page release performance enhancement as discussed in "9021-720 / Minidisk" on page 54. When these instructions are factored out, the number of DIAGNOSE instructions issued per command showed a 1.5% decrease from VM/ESA 1.0 to VM/ESA 1.1

RELEASE RUN ID	ESA 1.0 Y63F4809	ESA 1.1 Y64F480X
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**4800**	**4800**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
Response Time		
TRIV INT	0.052	0.049
NONTRIV INT	0.437	0.341
TOT INT	0.262	0.220
TOT INT ADJ	0.317	0.245
AVG FIRST (T)	0.495	0.330
AVG LAST (T)	0.745	0.507
Throughput		
AVG THINK (T)	25.56	25.51
ETR	203.80	189.27
ETR (T)	168.70	169.82
ETR RATIO	1.208	1.115
ITR (H)	189.12	201.64
ITR	38.12	37.49
EMUL ITR	64.38	59.13
ITRR (H)	1.000	1.066
ITRR	1.000	0.984
Proc. Usage		
PBT/CMD (H)	31.726	29.755
PBT/CMD	31.713	29.737
CP/CMD (H)	13.365	11.427
CP/CMD	12.922	10.894
EMUL/CMD (H)	18.355	18.322
EMUL/CMD	18.791	18.843
Processor Util.		
TOTAL (H)	535.21	505.32
TOTAL	535.00	505.00
UTIL/PROC (H)	89.20	84.22
UTIL/PROC	89.17	84.17
TVR(H)	1.73	1.62
TVR	1.69	1.58
Storage		
WKSET (V)	61	64
PGBLPGS	108K	108K
PGBLPGS/USER	23.0	23.0
FREEPGS	11511	12217
FREE UTIL	0.96	0.96
SHRPGS	1370	1322
Paging		
READS/SEC	246	298
WRITES/SEC	154	161
PAGE/CMD	2.371	2.703
XSTOR IN/SEC	1004	1086
XSTOR OUT/SEC	1207	1298
XSTOR/CMD	13.106	14.038
FAST CLR/CMD	7.789	5.618

RELEASE **RUN ID**	**ESA 1.0** **Y63F4809**	**ESA 1.1** **Y64F480X**
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**4800**	**4800**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
I/O		
VIO RATE	1151	1178
VIO/CMD	6.823	6.937
MDC READS	1014	1044
MDC WRITES	304	309
MDC MODS	239	241
MDC HIT RATIO	0.94	0.93
PRIVOPs		
PRIVOP/CMD	27.636	28.826
DIAG/CMD	15.111	21.217
DIAG 08/CMD	0.741	0.730
DIAG 10/CMD	5.305	0.012
DIAG 58/CMD	1.239	1.242
DIAG 98/CMD	0.314	0.312
DIAG A4/CMD	2.608	2.726
DIAG A8/CMD	1.689	1.684
DIAG 214/CMD	na	11.547
SIE/CMD	78.910	72.357
SIE INTCPT/CMD	54.448	50.650
FREE TOTL/CMD	133.540	96.476
VTAM Machines		
WKSET (V)	1529	1402
TOT CPU/CMD (V)	4.4557	3.9843
CP CPU/CMD (V)	2.5555	1.9921
VIRT CPU/CMD (V)	1.9002	1.9921
DIAG 98/CMD (V)	0.319	0.316
SFS Servers		
WKSET (V)	1101	997
TOT CPU/CMD (V)	4.3404	3.7171
CP CPU/CMD (V)	2.4271	1.7960
VIRT CPU/CMD (V)	1.9133	1.9211
FP REQ/CMD(Q)	1.337	1.338
IO/CMD (Q)	2.002	1.982
IO TIME/CMD (Q)	0.041	0.036
SFS TIME/CMD (Q)	0.127	0.088

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

Table 5. VM/ESA 1.0 and VM/ESA 1.1 SFS Regression on the 9021-720.

9021-580 / 35% SFS

The following is a description of the environment used to test VM 1.1 regression on a 9021-580.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

- PROCESSOR: 9021-580
- STORAGE:
  - RSTOR: 256M
  - XSTOR: 1G
- DASD:

	PACK NAME	TYPE
- SYSTEM:	RESPAK	3380-A
	SRVPAK	3380-A
	ESAP01	3380-A
	ESAOV1	3380-A
	ESAOV2	3380-A

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3390-A	2 - 3990-3	5	5	5	0	10
3380-A	1 - 3880-2	0	0	0	5	0

- TAPE: MONITOR 3480
- COMMUNICATIONS:

CTCA	NUMBER	CHANNEL SPEED
3088	1	4.5M

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2M
- USER CMS MODE: XA
- USER RELSHARE: 100
- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON
RWSERV1	SFS	32M/XA	1500	QUICKDSP ON
RWSERV2	SFS	32M/XA	1500	QUICKDSP ON
CRRSERV1	CRR	17M/XA	100	NONE

4) MEASUREMENT DISCUSSION

The following table summarizes the results of the regression measurements on the 9021-580 comparing VM/ESA 1.0 and VM/ESA 1.1 for the 35% SFS workload. These results were similar to the 9021-720 35% SFS regression results and the reader may want to refer to “9021-720 / 35% SFS” on page 58 for further details.

The 9021-580 experienced a 3.7% increase in internal throughput (ITR (H)) with almost all of this gain from a decrease of 10% in CP time per command (CP/CMD (H)). While the contributing factors to this improvement were the same as the 9021-720, there was less gain from the MDC Spin Lock Fix and Pending Page Release enhancements. The MDC Spin Lock Fix provides more benefit for systems with larger minidisk cache sizes and more processors while Pending Page Release reduces the number of SSKE and PTLB instructions which are more costly on systems with more processors. This system had only three processors and 1G of expanded storage while the 9021-720 had six processors and 2G of expanded storage.

A different network configuration was used for the 9021-580 runs. This system had only one VTAM machine with an internal VSCS server and a CTCA was used in place of 3745′s. This configuration, in particular the CTCA, resulted in lower external response times (AVG LAST (T)) for the 9021-580 runs than the 9021-720 runs and caused the RTM measured external throughput rate (ETR) to be smaller than the external measurement (ETR (T)). The CTCA allows transactions to come in so fast that CP is internally counting multiple transactions as a single transaction. This is due to the fact that CP has no good way of truly determining what is one complete transaction. As a result, it uses a time period where if a virtual machine is re-dispatched in this time period it counts the two dispatches as one transaction, even though they may actually be two distinct external transactions.

The 9021-580 had a similar increase in CP and CMS real storage requirements as the 9021-720.

The 9021-580 experienced a 40% reduction in external response time and a 38% reduction in internal response time (TOT INT ADJ). The dominant factor in the response time improvement was the SFS Checkpoint improvement in VM/ESA 1.1 (discussed in "9021-720 / 35% SFS"). Actually, the 9021-580 experienced a greater improvement in external response time than the 9021-720, which experienced a 33% improvement. Although the improvement in the SFS servers′ processor time per command (SFS TOT CPU/CMD (V)) was not as great on the 9021-580 (7%) as it was on the 9021-720 (14%), there was a 37% reduction in SFS I/O time per command (SFS IO TIME/CMD (Q)) and a 57% reduction in total SFS time per command (SFS TIME/CMD(Q)) for the 9021-580. These improvements were greater than the 9021-720, which saw 10% and 32% improvements respectively. Since the proportion of the internal response time due to SFS (SFS TIME/CMD (Q)/TOT INT ADJ) was the same on both systems (approximately 40%), the additional reduction in SFS time per command on the 9021-580 caused the better response time improvement.

The IUCV enhancements had less of an effect in this environment since an internal VSCS machine was used, eliminating the IUCV connections between VTAM and VSCS. It helped improve the VTAM servers′ processor time per command (VTAM TOT CPU/CMD (V)), which reduced external response times but this doesn′t appear to be the dominant cause of external response time improvement. Most of the external response time improvement actually came from improvements in internal response time (TOT INT ADJ). In fact, the VTAM servers′ processor time per command showed less of an improvement on the 9021-580 than it did on the 9021-720 and the delays in VTAM and the network overhead had less of an impact on external response times on the 9021-580 than it did on the 9021-720.

RELEASE RUN ID	ESA 1.0 Y33F2642	ESA 1.1 Y34F2644
Environment		
REAL STORAGE	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**
USERS	**2640**	**2640**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**3**	**3**
Response Time		
TRIV INT	0.110	0.108
NONTRIV INT	0.847	0.521
TOT INT	0.590	0.389
TOT INT ADJ	0.540	0.338
AVG FIRST (T)	0.320	0.210
AVG LAST (T)	0.650	0.387
Throughput		
AVG THINK (T)	25.39	25.46
ETR	86.82	82.71
ETR (T)	94.86	95.27
ETR RATIO	0.915	0.868
ITR (H)	104.91	108.78
ITR	31.99	31.54
EMUL ITR	50.63	47.63
ITRR (H)	1.000	1.037
ITRR	1.000	0.986
Proc. Usage		
PBT/CMD (H)	28.595	27.580
PBT/CMD	28.568	27.501
CP/CMD (H)	10.904	9.817
CP/CMD	10.436	9.237
EMUL/CMD (H)	17.687	17.759
EMUL/CMD	18.132	18.264
Processor Util.		
TOTAL (H)	271.26	262.75
TOTAL	271.00	262.00
UTIL/PROC (H)	90.42	87.58
UTIL/PROC	90.33	87.33
TVR(H)	1.62	1.55
TVR	1.58	1.51
Storage		
WKSET (V)	59	64
PGBLPGS	49861	49124
PGBLPGS/USER	18.9	18.6
FREEPGS	6210	6677
FREE UTIL	0.95	0.96
SHRPGS	872	1196
Paging		
READS/SEC	154	218
WRITES/SEC	81	77
PAGE/CMD	2.477	3.096
XSTOR IN/SEC	637	678
XSTOR OUT/SEC	743	778
XSTOR/CMD	14.548	15.283
FAST CLR/CMD	7.811	5.563

RELEASE RUN ID	ESA 1.0 Y33F2642	ESA 1.1 Y34F2644
Environment		
REAL STORAGE	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**
USERS	**2640**	**2640**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**3**	**3**
I/O		
VIO RATE	656	660
VIO/CMD	6.915	6.928
MDC READS	519	526
MDC WRITES	176	176
MDC MODS	119	111
MDC HIT RATIO	0.90	0.89
PRIVOPs		
PRIVOP/CMD	24.003	24.864
DIAG/CMD	15.888	22.320
DIAG 08/CMD	0.759	0.735
DIAG 10/CMD	5.313	0.010
DIAG 58/CMD	1.254	1.249
DIAG 98/CMD	0.590	0.756
DIAG A4/CMD	2.646	2.750
DIAG A8/CMD	1.729	1.627
DIAG 214/CMD	na	11.536
SIE/CMD	72.854	67.734
SIE INTCPT/CMD	48.812	45.382
FREE TOTL/CMD	118.742	95.371
VTAM Machines		
WKSET (V)	1120	925
TOT CPU/CMD (V)	3.3558	3.1023
CP CPU/CMD (V)	1.7569	1.4112
VIRT CPU/CMD (V)	1.5988	1.6911
DIAG 98/CMD (V)	0.597	0.759
SFS Servers		
WKSET (V)	1235	988
TOT CPU/CMD (V)	3.8711	3.6096
CP CPU/CMD (V)	1.9795	1.6211
VIRT CPU/CMD (V)	1.8916	1.9885
FP REQ/CMD(Q)	1.320	1.317
IO/CMD (Q)	1.821	1.803
IO TIME/CMD (Q)	0.052	0.033
SFS TIME/CMD (Q)	0.219	0.095
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM		

Table 6. VM/ESA 1.1 35% SFS Regression on 9021-580.

9121-480 / Minidisk

The following is a description of the environment used to test VM/ESA 1.1 minidisk regression on the 9121-480.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9121-480
- STORAGE:
   - RSTOR:       192M
   - XSTOR:        64M     All reserved for MDC

- DASD:
             PACK NAME   TYPE
   - SYSTEM: PSYS02      3380-A
             PSPT01      3380-A
             WKLD01      3380-A
             WKLD02      3380-A
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	11 - 3880-2	16	4	8	20	0
3380-A	2 - 3880-J23	0	0	0	4	0

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	2	44	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:    100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON, RESERVE 1050

4) MEASUREMENT DISCUSSION

The following table summarizes the results of the regression measurements on the 9121-480 comparing VM/ESA 1.0 and VM/ESA 1.1 for the minidisk workload. These results were similar to the 9021-720 minidisk regression results and the reader may want to refer to “9021-720 / Minidisk” on page 54 for further details.

The 9121-480 experienced a 3.2% increase in internal throughput ITR (H) with almost all of this gain from a decrease of 8.8% in CP time per command (CP/CMD (H)). External response times (AVG LAST (T)) improved 0.065 seconds, or 10.8%. Both the internal throughput and external response time show a smaller improvement than was experienced on the 9021-720. While the contributing factors to this improvement were the same as the 9021-720, as explained in “9021-580 / 35% SFS,” the MDC Spin Lock Fix and Pending Page Release improvement have less impact on smaller systems with fewer processors and less expanded storage for minidisk caching.

The 9121-480 had fewer users managed by the VSCS server. This resulted in a smaller portion of the external response time on the 9121-480 being due to VTAM and network delay than on the 9021-720. This can be seen by subtracting the internal response time (TOT INT ADJ) from the external response time. This difference was 0.565 (70% of total) for the VM/ESA 1.0 run on the 9021-720. For the 9121-480, it was only 0.253 (42% of total). The VTAM processor busy time per command (VTAM TOT CPU/CMD (V)) showed less improvement on the 9121-480 than on the 9021-720. With the smaller improvement on the 9121-480, and the fact that the VTAM and network delay was smaller on the 9121-480 to begin with, external response time did not improve as much as the 9021-720, which experienced a 35% reduction in external response time.

There was a similar increase in storage requirements for VM/ESA 1.1 on the 9121-480 as the 9021-720. There was a six page (8.8%) increase in average working set size (WKSET (V)) and there was an increase in the storage required by CP, shown as an increase in the FREEPGS and a decrease in the PGBLPGS. These increased storage requirements have a more negative effect on the 9121-480 since it is moderately storage constrained and has no expanded storage for paging. The effects of this can be seen by the approximate one page increase in the paging rate per command (PAGE/CMD). Most of this increase was in page reads (READS/SEC), which require the user virtual machine to wait for the paging operation to complete before it can continue. For the 9021-720, there was only a half page per command paging rate increase and a smaller percentage of the increase was from page reads. This increased paging also contributes to a smaller reduction in response time on the 9121-480 than on the 9021-720.

RELEASE RUN ID	ESA 1.0 L23R1770	ESA 1.1 L24R1770
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**1770**	**1770**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**2**	**2**
Response Time		
TRIV INT	0.068	0.065
NONTRIV INT	0.473	0.410
TOT INT	0.311	0.273
TOT INT ADJ	0.347	0.301
AVG FIRST (T)	0.385	0.355
AVG LAST (T)	0.600	0.535
Throughput		
AVG THINK (T)	25.71	25.74
ETR	70.17	69.14
ETR (T)	62.84	62.66
ETR RATIO	1.117	1.103
ITR (H)	69.53	71.78
ITR	38.80	39.62
EMUL ITR	62.51	61.15
ITRR (H)	1.000	1.032
ITRR	1.000	1.021
Proc. Usage		
PBT/CMD (H)	28.765	27.863
PBT/CMD	28.804	27.927
CP/CMD (H)	11.163	10.176
CP/CMD	10.981	9.894
EMUL/CMD (H)	17.597	17.681
EMUL/CMD	17.824	18.033
Processor Util.		
TOTAL (H)	180.76	174.59
TOTAL	181.00	175.00
UTIL/PROC (H)	90.38	87.30
UTIL/PROC	90.50	87.50
TVR(H)	1.63	1.58
TVR	1.62	1.55
Storage		
WKSET (V)	68	74
PGBLPGS	40079	39871
PGBLPGS/USER	22.6	22.5
FREEPGS	4415	4505
FREE UTIL	0.99	0.97
SHRPGS	877	891
Paging		
READS/SEC	466	507
WRITES/SEC	330	340
PAGE/CMD	12.668	13.517
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	8.037	5.458

RELEASE RUN ID	ESA 1.0 L23R1770	ESA 1.1 L24R1770
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**1770**	**1770**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**2**	**2**
I/O		
VIO RATE	533	541
VIO/CMD	8.482	8.634
MDC READS	346	359
MDC WRITES	167	169
MDC MODS	137	139
MDC HIT RATIO	0.92	0.92
PRIVOPs		
PRIVOP/CMD	20.903	21.869
DIAG/CMD	18.075	25.202
DIAG 08/CMD	0.764	0.766
DIAG 10/CMD	5.649	0.016
DIAG 58/CMD	1.257	1.245
DIAG 98/CMD	0.446	0.463
DIAG A4/CMD	3.867	4.038
DIAG A8/CMD	2.085	2.043
DIAG 214/CMD	na	12.320
SIE/CMD	65.693	58.536
SIE INTCPT/CMD	45.328	40.975
FREE TOTL/CMD	106.448	84.309
VTAM Machines		
WKSET (V)	1195	1217
TOT CPU/CMD (V)	4.9333	4.4861
CP CPU/CMD (V)	2.7849	2.2253
VIRT CPU/CMD (V)	2.1484	2.2608
DIAG 98/CMD (V)	0.457	0.466
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 7. VM/ESA 1.1 Minidisk Regression on the 9121-480.

9121-480 / 35% SFS

The following is a description of the environment used to test VM/ESA 1.1 regression on the 9121-480.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

- PROCESSOR: 9121-480
- STORAGE:
  - RSTOR: 192M
  - XSTOR: 64M All reserved for MDC

- DASD:

	PACK NAME	TYPE
- SYSTEM:	PSYS02	3380-A
	PSPT01	3380-A
	WKLD01	3380-A
	WKLD02	3380-A

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-A	11 - 3880-2	16	4	8	0	20
3380-A	2 - 3880-J23	0	0	0	4	0

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	2	44	56Kb

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2M
- USER CMS MODE: XA
- USER RELSHARE: 100

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON, RESERVE 850
RWSERV1	SFS	32M/XA	1500	QUICKDSP ON, RESERVE 1300
RWSERV2	SFS	32M/XA	1500	QUICKDSP ON, RESERVE 1300
CRRSERV1	CRR	17M/XA	100	NONE

4) MEASUREMENT DISCUSSION

The following table summarizes the results of the regression measurements on the 9121-480 comparing VM/ESA 1.0 and VM/ESA 1.1 for the 35% SFS workload. These results were similar to the 9021-720 and 9121-580 35% SFS regression results and the reader may want to refer to “9021-720 / 35% SFS” on page 58 and “9021-580 / 35% SFS” on page 62 for further details.

The 9121-480 experienced a 4.5% increase in internal throughput (ITR (H)) with almost all of this gain from a decrease of 11% in CP time per command (CP/CMD (H)). External response times (AVG LAST (T)) on the 9121-480 were reduced by 0.275 seconds (29%), compared to a reduction of 0.25 seconds (33%) on the 9021-720. While the contributing factors to this improvement were the same as for the 9021-720, as discussed in “9021-580 / 35% SFS,” the MDC Spin Lock Fix and Pending Page Release enhancement provide less benefit for smaller systems with fewer processors and less expanded storage for minidisk caching.

The amount of internal and external response time improvement was greater than the corresponding minidisk only comparison, mainly due to the SFS Checkpoint improvement. Additionally, the IUCV enhancements provided more of an improvement for the 35% SFS workload than it did the minidisk workload, due to the use of APPC/VM and *BLOCKIO by SFS.

The 9121-480 35% SFS workload experienced a similar increase in CP and CMS real storage requirements as the 9121-480 minidisk workload. This resulted in a similar increase in relatively expensive page reads (READS/SEC) which contributed to a slightly smaller reduction in response times on the 9121-480 than on the 9021-720. See “9121-480 / Minidisk” for further details.

The proportion of the external response time belonging to the VTAM and network delay was smaller for the 9121-480 than it was the 9021-720. This resulted in a smaller improvement in external response time from the reduction in VTAM processing time by the IUCV improvements.

The improvement in the SFS servers′ processor time per command (SFS TOT CPU/CMD (V)) was not as great on the 9121-480 (7.2%) as it was on the 9021-720 (14.0%). SFS time per command was reduced 0.072 seconds (32%) on the 9121-480, which was similar to the 9021-720 that had a reduction of 0.039 seconds (33%). The proportion of the internal response time (TOT INT ADJ) was 35% for the 9121-480, (compared to 40% for the 9021-720), causing the 9121-480 to benefit slightly less from the SFS Checkpoint improvements than the 9021-720.

RELEASE RUN ID	ESA 1.0 L23F1484	ESA 1.1 L24F1480
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**1480**	**1480**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**2**	**2**
Response Time		
TRIV INT	0.071	0.069
NONTRIV INT	0.895	0.639
TOT INT	0.564	0.424
TOT INT ADJ	0.636	0.460
AVG FIRST (T)	0.495	0.380
AVG LAST (T)	0.945	0.670
Throughput		
AVG THINK (T)	25.52	25.34
ETR	58.77	57.23
ETR (T)	52.10	52.78
ETR RATIO	1.128	1.084
ITR (H)	58.55	61.16
ITR	33.01	33.19
EMUL ITR	55.46	52.81
ITRR (H)	1.000	1.045
ITRR	1.000	1.005
Proc. Usage		
PBT/CMD (H)	34.158	32.699
PBT/CMD	34.168	32.777
CP/CMD (H)	14.083	12.523
CP/CMD	13.821	12.315
EMUL/CMD (H)	20.068	20.167
EMUL/CMD	20.347	20.462
Processor Util.		
TOTAL (H)	177.95	172.59
TOTAL	178.00	173.00
UTIL/PROC (H)	88.97	86.29
UTIL/PROC	89.00	86.50
TVR(H)	1.70	1.62
TVR	1.68	1.60
Storage		
WKSET (V)	73	81
PGBLPGS	41095	40807
PGBLPGS/USER	27.8	27.6
FREEPGS	3748	3931
FREE UTIL	0.94	0.93
SHRPGS	1067	1082
Paging		
READS/SEC	403	462
WRITES/SEC	318	325
PAGE/CMD	13.840	14.911
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	7.832	5.646

RELEASE RUN ID	ESA 1.0 L23F1484	ESA 1.1 L24F1480
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**1480**	**1480**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**2**	**2**
I/O		
VIO RATE	357	363
VIO/CMD	6.853	6.878
MDC READS	278	290
MDC WRITES	96	100
MDC MODS	56	58
MDC HIT RATIO	0.87	0.87
PRIVOPs		
PRIVOP/CMD	28.974	30.154
DIAG/CMD	15.849	22.385
DIAG 08/CMD	0.749	0.758
DIAG 10/CMD	5.183	0.000
DIAG 58/CMD	1.248	1.250
DIAG 98/CMD	0.518	0.530
DIAG A4/CMD	2.495	2.652
DIAG A8/CMD	1.843	1.705
DIAG 214/CMD	na	11.406
SIE/CMD	77.626	71.049
SIE INTCPT/CMD	55.115	51.866
FREE TOTL/CMD	133.485	96.210
VTAM Machines		
WKSET (V)	1072	1035
TOT CPU/CMD (V)	5.1934	4.6313
CP CPU/CMD (V)	2.9646	2.3262
VIRT CPU/CMD (V)	2.2288	2.3051
DIAG 98/CMD (V)	0.536	0.535
SFS Servers		
WKSET (V)	1364	1355
TOT CPU/CMD (V)	4.4256	4.1051
CP CPU/CMD (V)	2.2715	1.9367
VIRT CPU/CMD (V)	2.1542	2.1683
FP REQ/CMD(Q)	1.328	1.333
IO/CMD (Q)	1.825	1.823
IO TIME/CMD (Q)	0.071	0.055
SFS TIME/CMD (Q)	0.224	0.152

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

Table 8. VM/ESA 1.1 35% SFS Regression on 9121-480.

9121-320 / Minidisk

The following is a description of the environment used to test VM/ESA 1.1 minidisk regression on the 9121-320.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9121-320
- STORAGE:
   - RSTOR:       192M
   - XSTOR:        64M     All reserved for MDC

- DASD:
             PACK NAME   TYPE
   - SYSTEM: PSYS02      3380-A
             PSPT01      3380-A
             WKLD01      3380-A
             WKLD02      3380-A
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	11 - 3880-2	16	4	8	10	0
3380-A	2 - 3880-J23	0	0	0	4	0

```
- TAPE:        MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	2	44	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:   100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON

4) MEASUREMENT DISCUSSION

The following table summarizes the results of the regression measurements on the 9121-320 comparing VM/ESA 1.0 and VM/ESA 1.1 for the minidisk workload. These results were similar to the 9021-720 and 9121-480 minidisk regression results. See "9021-720 / Minidisk" on page 54 and "9121-480 / Minidisk" on page 66 for more details.

The 9121-320 experienced a 4.5% increase in internal throughput (ITR (H)) with almost all of this gain from a decrease of 10.3% in CP time per command (CP/CMD (H)). External response times (AVG LAST (T)) improved 0.047 seconds, or 7.6%. This is a smaller improvement than was shown on the 9121-480 and the 9021-720. Since the 9121-320 is a uniprocessor, the MDC Spin Lock Fix provided no ITR improvements. As explained in "9021-580 / 35% SFS," the Pending Page Release enhancement has less effect on systems with fewer processors.

Although the VTAM processor busy time per command showed a similar improvement (VTAM TOT CPU/CMD (V)) for the 9121-320 as it did on the 9121-480, the VTAM IUCV enhancements provided even less benefit in terms of external response time on the 9121-320 than on the 9121-480. There were fewer users for the VSCS server to manage on the 9121-320 which resulted in a smaller percentage of the external response time due to the VTAM and network delay. In this case, the VTAM and network delay in the external response times was 0.161 seconds (30%), compared to 0.234 (42%) for the 9121-480.

There was a similar increase in real storage requirements for the 9121-320 as there was on the 9121-480. The increased storage requirements resulted in an increase of a little over one page per command in the paging rate (PAGE/CMD). This was slightly less of an increase than on the 9121-480 since it is not as storage constrained. Again, since there was no expanded storage for paging on the 9121-320, this was a direct increase in page I/Os to the DPA, which resulted in a smaller improvement in external response times than the 9021-720 experienced.

RELEASE RUN ID	ESA 1.0 L13R0911	ESA 1.1 L14R0910
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**910**	**910**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**1**	**1**
Response Time		
TRIV INT	0.072	0.068
NONTRIV INT	0.521	0.453
TOT INT	0.365	0.319
TOT INT ADJ	0.376	0.329
AVG FIRST (T)	0.290	0.280
AVG LAST (T)	0.530	0.490
Throughput		
AVG THINK (T)	26.02	25.82
ETR	33.08	33.39
ETR (T)	32.15	32.41
ETR RATIO	1.029	1.030
ITR (H)	35.36	36.94
ITR	36.39	38.09
EMUL ITR	58.53	58.55
ITRR (H)	1.000	1.045
ITRR	1.000	1.047
Proc. Usage		
PBT/CMD (H)	28.282	27.070
PBT/CMD	28.301	27.148
CP/CMD (H)	10.908	9.784
CP/CMD	10.574	9.564
EMUL/CMD (H)	17.366	17.279
EMUL/CMD	17.727	17.585
Processor Util.		
TOTAL (H)	90.94	87.75
TOTAL	91.00	88.00
UTIL/PROC (H)	90.94	87.75
UTIL/PROC	91.00	88.00
TVR(H)	1.63	1.57
TVR	1.60	1.54
Storage		
WKSET (V)	70	81
PGBLPGS	43513	43242
PGBLPGS/USER	47.8	47.5
FREEPGS	2293	2373
FREE UTIL	0.93	0.93
SHRPGS	792	784
Paging		
READS/SEC	224	252
WRITES/SEC	149	159
PAGE/CMD	11.600	12.679
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	8.117	5.399

RELEASE RUN ID	ESA 1.0 L13R0911	ESA 1.1 L14R0910
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**910**	**910**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**1**	**1**
I/O		
VIO RATE	275	279
VIO/CMD	8.553	8.607
MDC READS	178	187
MDC WRITES	85	88
MDC MODS	70	72
MDC HIT RATIO	0.92	0.92
PRIVOPs		
PRIVOP/CMD	23.413	24.160
DIAG/CMD	18.289	25.339
DIAG 08/CMD	0.778	0.771
DIAG 10/CMD	5.691	0.000
DIAG 58/CMD	1.244	1.234
DIAG 98/CMD	0.435	0.432
DIAG A4/CMD	3.856	4.041
DIAG A8/CMD	2.115	2.005
DIAG 214/CMD	na	12.340
SIE/CMD	69.292	61.485
SIE INTCPT/CMD	49.890	44.884
FREE TOTL/CMD	114.667	86.720
VTAM Machines		
WKSET (V)	753	1055
TOT CPU/CMD (V)	5.4598	4.8675
CP CPU/CMD (V)	3.1273	2.4337
VIRT CPU/CMD (V)	2.3325	2.4337
DIAG 98/CMD (V)	0.457	0.461
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 9. VM/ESA 1.1 Minidisk Regression on 9121-320.

9121-320 / 35% SFS

The following is a description of the environment used to test VM/ESA 1.1 regression on the 9121-320.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9121-320
- STORAGE:
   - RSTOR:       192M
   - XSTOR:        64M     All reserved for MDC

- DASD:
            PACK NAME   TYPE
   - SYSTEM: PSYS02     3380-A
             PSPT01     3380-A
             WKLD01     3380-A
             WKLD02     3380-A
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	11 - 3880-2	16	4	8	0	10
3380-A	2 - 3880-J23	0	0	0	4	0

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	1	44	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:    100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
RWSERV1	SFS	32M/XA	1500	QUICKDSP ON
CRRSERV1	CRR	17M/XA	100	NONE

4) MEASUREMENT DISCUSSION

The following table summarizes the results of the regression measurements on the 9121-320 comparing VM/ESA 1.0 and VM/ESA 1.1 for the 35% SFS workload. These results were similar to 9021-720 and 9121-480 35% SFS regression results. The reader may want to refer to “9021-720 / 35% SFS” on page 58 and “9121-480 / 35% SFS” on page 70 for further details.

The 9121-320 experienced a 4.9% increase in internal throughput (ITR (H)) with almost all of this gain from a decrease of 11.6% in CP time per command (CP/CMD (H)). External response times (AVG LAST (T)) on the 9121-320 were reduced by 0.26 seconds (29%), compared to a reduction of 0.25 seconds (33%) on the 9021-720. Since the 9121-320 is a uniprocessor, the MDC Spin Lock fix provided no ITR improvements. As explained in “9021-580 / 35% SFS,” the Pending Page Release enhancement has less effect on systems with fewer processors.

The amount of internal and external response time improvement was greater than the amount in the corresponding minidisk only comparison, due to the SFS Checkpoint Improvement. This was responsible for most of the 36% reduction in total SFS time per command (SFS TIME/CMD(Q)) for the 9121-320. Additionally, the IUCV enhancements provided more of an improvement for the 35% SFS workload than it did the minidisk workload, due to the use of APPC/VM and *BLOCKIO by SFS.

There was a similar increase in real storage requirements for the 9121-320 as there was on the 9121-480. The increased storage requirements resulted in an increase of 0.8 pages per command in the paging rate (PAGE/CMD). This was slightly less of an increase than on the 9121-480 since it was not as storage constrained. Again, since there was no expanded storage for paging on the 9121-320, this was a direct increase in page I/Os to the DPA, which resulted in less improvement in external response times than the 9021-720 experienced.

Although the VTAM processor busy time per command showed a similar improvement (VTAM TOT CPU/CMD (V)) for the 9121-320 as it did on the 9121-480, the VTAM IUCV enhancements provided even less benefit in terms of external response time on the 9121-320 than on the 9121-480. There were fewer users for the VSCS server to manage on the 9121-320 which resulted in a slightly smaller percentage of the external response time due to the VTAM and network delay.

RELEASE RUN ID	ESA 1.0 L13F0771	ESA 1.1 L14F0770
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**770**	**770**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**1**	**1**
Response Time		
TRIV INT	0.072	0.069
NONTRIV INT	1.005	0.656
TOT INT	0.682	0.452
TOT INT ADJ	0.712	0.469
AVG FIRST (T)	0.380	0.330
AVG LAST (T)	0.900	0.640
Throughput		
AVG THINK (T)	25.38	25.37
ETR	28.38	28.59
ETR (T)	27.19	27.58
ETR RATIO	1.044	1.037
ITR (H)	30.09	31.56
ITR	31.38	32.73
EMUL ITR	52.18	51.60
ITRR (H)	1.000	1.049
ITRR	1.000	1.043
Proc. Usage		
PBT/CMD (H)	33.233	31.688
PBT/CMD	33.095	31.546
CP/CMD (H)	13.466	11.900
CP/CMD	13.238	11.603
EMUL/CMD (H)	19.757	19.777
EMUL/CMD	19.857	19.943
Processor Util.		
TOTAL (H)	90.38	87.39
TOTAL	90.00	87.00
UTIL/PROC (H)	90.38	87.39
UTIL/PROC	90.00	87.00
TVR(H)	1.68	1.60
TVR	1.67	1.58
Storage		
WKSET (V)	75	88
PGBLPGS	43975	43749
PGBLPGS/USER	57.1	56.8
FREEPGS	2001	2117
FREE UTIL	0.92	0.92
SHRPGS	1000	1038
Paging		
READS/SEC	202	224
WRITES/SEC	143	147
PAGE/CMD	12.686	13.452
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	7.832	5.584

RELEASE RUN ID	ESA 1.0 L13F0771	ESA 1.1 L14F0770
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**770**	**770**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**1**	**1**
I/O		
VIO RATE	187	193
VIO/CMD	6.876	6.998
MDC READS	153	162
MDC WRITES	51	52
MDC MODS	37	37
MDC HIT RATIO	0.91	0.91
PRIVOPs		
PRIVOP/CMD	31.462	32.175
DIAG/CMD	16.026	22.826
DIAG 08/CMD	0.735	0.725
DIAG 10/CMD	5.258	0.000
DIAG 58/CMD	1.250	1.233
DIAG 98/CMD	0.478	0.508
DIAG A4/CMD	2.537	2.719
DIAG A8/CMD	1.802	1.740
DIAG 214/CMD	na	11.567
SIE/CMD	80.273	73.680
SIE INTCPT/CMD	59.402	55.997
FREE TOTL/CMD	142.343	98.953
VTAM Machines		
WKSET (V)	561	542
TOT CPU/CMD (V)	5.5770	4.9152
CP CPU/CMD (V)	3.2277	2.4979
VIRT CPU/CMD (V)	2.3493	2.4173
DIAG 98/CMD (V)	0.514	0.529
SFS Servers		
WKSET (V)	1505	1463
TOT CPU/CMD (V)	4.2287	3.8476
CP CPU/CMD (V)	2.2063	1.8331
VIRT CPU/CMD (V)	2.0224	2.0144
FP REQ/CMD(Q)	1.343	1.357
IO/CMD (Q)	1.923	1.888
IO TIME/CMD (Q)	0.038	0.031
SFS TIME/CMD (Q)	0.149	0.096
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM		

Table 10. VM/ESA 1.1 35% SFS Regression on 9121-320.

9221-170 / Minidisk

The Following is a description of the environment used for the minidisk regression measurements on the 9221-170.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9221-170
- STORAGE:
   - RSTOR:   48M
   - XSTOR:   16M (all reserved for MDC)

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  H3AP01     3380
              H3SRV      3380
              H3RES      3380
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-A	3 - 3380-03	2	2	2	4	0
3380-D	1 - 3380-03	1	0	1	1	0

(NUMBER OF PACKS: PAGE, SPOOL, TDISK, USER, SERVER)

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINESPEED
3088-02	1	4.5M

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR:  BACTRIAN
- CMS BLOCKSIZE:     4K
- USER VM SIZE:      2M
- USER CMS MODE:     370
- USER RELSHARE:     100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON

4) MEASUREMENT DISCUSSION

The following table summarizes the results of the regression measurements on the 9221-170 comparing VM/ESA 1.0 and VM/ESA 1.1 for the minidisk-only CMS intensive environment. The same system tuning is done for both measurements except for IPOLL ON, which is new in VM/ESA 1.1 See section "GCS IPOLL Option" on page 40.

The 9221-170 experienced a 4.3% increase in internal throughput (ITR (H)) with almost all of this gain from a decrease of 7.6% in CP time per command (CP/CMD (H)). This drop in CP CPU time is due largely to CP Fast Dynamic Linkage and IUCV enhancements. Since the 9221-170 is a uniprocessor, the MDC Spin Lock Fix provided no ITR improvements. As explained in "9021-580 / 35% SFS" on page 62, the Pending Page Release enhancement had less effect on systems with fewer processors. External response time (AVG LAST (T)) decreased by 0.12 seconds (17.9%) and internal response time (TOT INT ADJ) decreased by 0.12 seconds (18.5%). Therefore, all of the improvement in external response time was due to the improvement in internal response time.

Real storage requirements increased somewhat for VM/ESA 1.1 as evidenced by the overall increase in paging (PAGE/CMD); however, migration from VM/ESA 1.0 to VM/ESA 1.1 on the 9021-720, 9121-480, and 9121-320, showed a decrease in the PGBLPGS. This is not true for the 9221-170 because control blocks in VM/ESA 1.0 were created when excess expanded storage was attached to an idle user. Other systems may not experience this since there is an APAR (VM45743) available to correct this problem. This APAR was not applied to VM/ESA 1.0 on the 9221-170. There was also a growth in CP storage required for CP control block as evidenced by the increase in FREEPGS. The increased storage requirements affected the 9221-170 more than the 9021 and 9121 since I/O is handled less efficiently on the 9221.

RELEASE RUN ID	ESA 1.0 H13R0280	ESA 1.1 H14R0287
Environment		
REAL STORAGE	**48M**	**48M**
EXP. STORAGE	**16M**	**16M**
USERS	**280**	**280**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
Response Time		
TRIV INT	0.181	0.172
NONTRIV INT	1.021	0.810
TOT INT	0.776	0.624
TOT INT ADJ	0.650	0.529
AVG FIRST (T)	0.290	0.270
AVG LAST (T)	0.670	0.550
Throughput		
AVG THINK (T)	28.58	28.39
ETR	8.06	8.23
ETR (T)	9.63	9.71
ETR RATIO	0.837	0.848
ITR (H)	11.41	11.90
ITR	9.55	10.10
EMUL ITR	15.64	15.99
ITRR (H)	1.000	1.043
ITRR	1.000	1.058
Proc Usage		
PBT/CMD (H)	87.646	84.060
PBT/CMD	88.274	84.486
CP/CMD (H)	39.357	36.380
CP/CMD	34.271	30.910
EMUL/CMD (H)	48.289	47.680
EMUL/CMD	54.003	53.577
Processor Util		
TOTAL (H)	84.40	81.59
TOTAL	85.00	82.00
UTIL/PROC (H)	84.40	81.59
UTIL/PROC	85.00	82.00
TVR(H)	1.82	1.76
TVR	1.63	1.58
Storage		
WKSET (V)	75	77
PGBLPGS	9359	9520
PGBLPGS/USER	33.4	34.0
FREEPGS	767	804
FREE UTIL	0.88	0.88
SHRPGS	706	900
Paging		
READS/SEC	56	62
WRITES/SEC	46	49
PAGE/CMD	10.593	11.437
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	8.827	5.461

RELEASE RUN ID	ESA 1.0 H13R0280	ESA 1.1 H14R0287
Environment		
REAL STORAGE	**48M**	**48M**
EXP. STORAGE	**16M**	**16M**
USERS	**280**	**280**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
I/O		
VIO RATE	82	82
VIO/CMD	8.516	8.449
MDC READS	53	53
MDC WRITES	25	26
MDC MODS	20	21
MDC HIT RATIO	0.91	0.91
PRIVOPs		
PRIVOP/CMD	16.685	14.426
DIAG/CMD	20.652	26.002
DIAG 08/CMD	0.623	0.618
DIAG 10/CMD	5.920	0.000
DIAG 58/CMD	1.142	1.236
DIAG 98/CMD	2.492	2.473
DIAG A4/CMD	3.946	3.812
DIAG A8/CMD	1.973	2.061
DIAG 214/CMD	na	11.334
SIE/CMD	69.061	56.977
SIE INTCPT/CMD	50.415	42.163
FREE TOTL/CMD	109.148	97.571
VTAM Machines		
WKSET (V)	218	211
TOT CPU/CMD (V)	19.8608	18.3362
CP CPU/CMD (V)	9.7397	8.1785
VIRT CPU/CMD (V)	10.1211	10.1577
DIAG 98/CMD (V)	2.515	2.519
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 11. 9221-170 Minidisk Comparing VM/ESA 1.1 to VM/ESA 1.0

CMS Intensive Migration from VM/ESA 1.0 370 Feature on 9221

9221-170 / Minidisk

The following is a description of the environment used for the minidisk regression measurements on the 9221-170.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9221-170
- STORAGE:
   CP/370  - RSTOR:   64M
           - XSTOR:    0M
   ESA
           - RSTOR:   48M,64M,240M (see table)
           - XSTOR:    0M,16M      (see table), (all reserved for MDC)

- DASD:
            PACK NAME  TYPE
   - SYSTEM: PRF05     3380
     CP/370  PRF01     3380
             PRFRES    3380

   - SYSTEM: H3AP01    3380
     ESA     H3SRV     3380
             H3RES     3380
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
		NUMBER OF PACKS				
3380-A	3 - 3380-03	2	2	2	4	0
3380-D	1 - 3380-03	1	0	1	1	0

```
- TAPE:        MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINESPEED
3088-02	1	4.5M

3) SOFTWARE CONFIGURATION

```
- DRIVER:             TPNS
- THINK TIME DISTR:   BACTRIAN
- CMS BLOCKSIZE:      4K
- USER VM SIZE:       2M
- USER CMS MODE:      370
- USER PRIORITY:       64  CP/370
- USER RELSHARE:      100  ESA
- SERVER MACHINES:
```

	SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
CP/370					
	VTAM	VTAM/VSCS	16M/370		PRIORITY 1, QDROP OFF USERS, FAVOR 100, FAVOR
ESA					
	VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON

4) MEASUREMENT DISCUSSION

The following table summarizes the results of the regression measurements on the 9221-170 comparing VM/ESA 1.0 370 Feature to VM/ESA 1.1 for the minidisk-only CMS intensive workload.

When migrating from VM/ESA 1.0 370 Feature (H17R0281) to VM/ESA 1.1 without tuning (H14R0283), internal throughput (ITR(H)) decreased by 13.7%, external response time (AVG LAST (T)) increased by 11.3%, and PBT/CMD increased by 15.9%. VTAM TOT/CPU/CMD increased by 74.8% and VTAM VIRT CPU/CMD increased by 60.3% due to an increase in CTC I/Os.

The VTAM working set decreased from 529 pages in VM/ESA 1.0 370 Feature to 207 pages in VM/ESA 1.1 measurement without tuning. The apparent change results from the method of accounting for the GCS and VTAM shared pages. In VM/ESA 1.0 370 Feature, these pages are counted as part of the working set for the VTAM virtual machine. In VM/ESA 1.1, these pages are counted as part of the shared pages (SHRPGS), which increased from 286 to 785 between these two measurements.

Tuning strategies were implemented for VM/ESA 1.1 See "Recommended 9221 Tuning" on page 223 for a comparison of the untuned measurement (H14R0283) to the tuned measurement (H14R0287).

Comparing the tuned VM/ESA 1.1 measurement (H14R0287) back to VM/ESA 1.0 370 Feature (H17R0281) showed the internal throughput gap closed to 5.5% and external response time improved by 31.3%.

Increasing storage by 192M in VM/ESA 1.1 improved internal throughput by 6.6% and external response time by 21.8%. Paging (PAGE/CMD) went to zero. Working sets (WRKSET(V)) are artificially inflated since no pages are ever stolen from the excess real storage. Look at runs H14R0287 and H14R0286.

Note: 9221 processors configured with integrated I/O controllers, running in ESA/390, mode can only use 128 MB of main storage. If the installed processor storage on these machines is greater than 128 MB, the remaining storage may be used for expanded storage. All of the runs shown here used only channel-attached devices, making the use of more than 128 MB of main storage possible.

Comparing the tuned VM/ESA 1.1 with extra storage measurement (H14R0286) back to VM/ESA 1.0 370 Feature (H17R0281) shows internal throughput improved 0.7% and external response time improved by 46.3%.

An additional measurement not shown demonstrates that increasing storage by 64M in VM/ESA 1.1, instead of 192M, improved internal throughput only 1.0%.

9221 TUNING RELEASE RUN ID	N/A ESA 1.0 (370) H17R0281	NO ESA 1.1 H14R0283	YES ESA 1.1 H14R0287	YES ESA 1.1 H14R0286
Environment				
REAL STORAGE	**64M**	**64M**	**48M**	**240M**
EXP. STORAGE	**0M**	**0M**	**16M**	**16M**
USERS	**280**	**280**	**280**	**280**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**	**0**
PROCESSORS	**1**	**1**	**1**	**1**
Response Time				
TRIV INT	0.180	0.226	0.172	0.124
NONTRIV INT	3.140	1.383	0.810	0.636
TOT INT	0.500	1.061	0.624	0.486
TOT INT ADJ	0.461	0.869	0.529	0.415
AVG FIRST (T)	0.460	0.320	0.270	0.210
AVG LAST (T)	0.800	0.890	0.550	0.430
Throughput				
AVG THINK (T)	28.57	28.41	28.39	28.48
ETR	8.79	7.89	8.23	8.28
ETR (T)	9.53	9.63	9.71	9.69
ETR RATIO	0.922	0.819	0.848	0.854
ITR (H)	12.59	10.86	11.90	12.68
ITR	11.62	8.91	10.10	10.83
EMUL ITR	19.71	14.48	15.99	16.50
ITRR (H)	1.000	0.863	0.945	1.007
ITRR	1.000	0.766	0.869	0.932
Proc. Usage				
PBT/CMD (H)	79.416	92.057	84.060	78.848
PBT/CMD	79.297	92.420	84.486	79.447
CP/CMD (H)	34.756	42.339	36.380	31.782
CP/CMD	32.533	35.306	30.910	27.858
EMUL/CMD (H)	44.659	49.718	47.680	47.066
EMUL/CMD	46.765	57.113	53.577	51.589
Processor Util.				
TOTAL (H)	75.72	88.65	81.59	76.42
TOTAL	75.61	89.00	82.00	77.00
UTIL/PROC (H)	75.72	88.65	81.59	76.42
UTIL/PROC	75.61	89.00	82.00	77.00
TVR(H)	1.78	1.85	1.76	1.68
TVR	1.70	1.62	1.58	1.54
Storage				
WKSET (V)	90	76	77	136
PGBLPGS	15006	13585	9520	57691
PGBLPGS/USER	53.6	48.5	34.0	206.0
FREEPGS	864	804	804	786
FREE UTIL	na	0.88	0.88	0.88
SHRPGS	286	785	900	3522
Paging				
READS/SEC	43	83	62	0
WRITES/SEC	23	48	49	0
PAGE/CMD	6.857	13.603	11.437	0.000
XSTOR IN/SEC	na	0	0	0
XSTOR OUT/SEC	na	0	0	0
XSTOR/CMD	na	0.000	0.000	0.000
FAST CLR/CMD	na	5.711	5.461	5.365

9221 TUNING RELEASE RUN ID	N/A ESA 1.0 (370) H17R0281	NO ESA 1.1 H14R0283	YES ESA 1.1 H14R0287	YES ESA 1.1 H14R0286
Environment				
REAL STORAGE	**64M**	**64M**	**48M**	**240M**
EXP. STORAGE	**0M**	**0M**	**16M**	**16M**
USERS	**280**	**280**	**280**	**280**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**	**0**
PROCESSORS	**1**	**1**	**1**	**1**
I/O				
VIO RATE	144	80	82	81
VIO/CMD	15.145	8.307	8.449	8.357
MDSK/CMD	5.558	na	na	na
MDC READS	na	0	53	53
MDC WRITES	na	0	26	26
MDC MODS	na	0	21	21
MDC HIT RATIO	na	0.00	0.91	0.91
PRIVOPs				
PRIVOP/CMD (R)	9.077	18.028	14.426	14.368
DIAG/CMD (R)	16.375	26.436	26.002	25.900
DIAG 08/CMD	na	0.623	0.618	0.619
DIAG 10/CMD	na	0.000	0.000	0.000
DIAG 58/CMD	na	1.142	1.236	1.238
DIAG 98/CMD	na	2.908	2.473	2.373
DIAG A4/CMD	na	3.946	3.812	3.921
DIAG A8/CMD	na	1.765	2.061	1.960
DIAG 214/CMD	na	11.423	11.334	11.350
SIE/CMD	na	70.301	56.977	53.756
SIE INTCPT/CMD	na	47.102	42.163	41.930
FREE TOTL/CMD	na	101.350	97.571	96.988
VTAM Machines				
WKSET (V)	529	207	211	1002
TOT CPU/CMD (V)	11.8812	20.7685	18.3362	17.8085
CP CPU/CMD (V)	4.7818	9.3869	8.1785	7.9570
VIRT CPU/CMD (V)	7.0994	11.3816	10.1577	9.8515
DIAG 98/CMD (V)	1.834	3.004	2.519	2.470

Note: T=TPNS, V=VMPRF (ESA) or VMMAP (370), H=Hardware Monitor, R=RTM, Unmarked=RTM (ESA) or VMMAP (370)

Table 12. 9221-170 Minidisk Comparing VM/ESA 1.0 370 Feature and VM/ESA 1.1

9221-170 / 35% SFS

The following is a description of the environment used for the 35% SFS measurements on the 9221-170.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9221-170
- STORAGE:
   CP/370 - RSTOR:   64M
          - XSTOR:   0M
   ESA
          - RSTOR:   48M
          - XSTOR:   16M (all reserved for MDC)

- DASD:
             PACK NAME  TYPE
   - SYSTEM: PRF05      3380
     CP/370  PRF01      3380
             PRFRES     3380

   - SYSTEM: H3AP01     3380
     ESA     H3SRV      3380
             H3RES      3380
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	3 - 3380-03	2	2	2	4	5
3380-D	1 - 3380-03	1	0	1	1	1

```
- TAPE:          MONITOR  3480

- COMMUNICATIONS:
                      CONTROLLER  NUMBER  LINESPEED
                      3088-02       1       4.5M
```

3) SOFTWARE CONFIGURATION

```
- DRIVER:              TPNS
- THINK TIME DISTR:    BACTRIAN
- CMS BLOCKSIZE:       4K
- USER VM SIZE:        2M
- USER CMS MODE:       370
- USER RELSHARE:        64  CP/370
- USER RELSHARE:       100  ESA
- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
CP/370				
VTAM	VTAM/VSCS	16M/370		PRIORITY 1, QDROP OFF USERS, FAVOR 100, FAVOR FAVOR 100, FAVOR
RWSERVE1 (R/W)	SFS	16M/370		PRIORITY 1, QDROP OFF USERS
CRRRECOV	CRR	16M/370		NONE
ESA				
VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON
RWSERVE1 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
CRRRECOV	CRR	32M/XC	100	NONE

4) MEASUREMENT DISCUSSION

This section summarizes the results of SFS measurements when migrating from VM/ESA 1.0 370 Feature to VM/ESA 1.1 with tuning.

Internal throughput (ITR(H)) decreased by 7.5%. External response time (AVG LAST (T)) improved by 12.3% even though paging (PAGE/CMD) increased. The response time improvement was due to the fact that VM/ESA 1.1′s minidisk caching eliminated some of the DASD I/Os. See section "Recommended 9221 Tuning" on page 223 for more details. The same trends were experienced when migrating in the minidisk environment. For more details, see section "9221-170 / Minidisk" on page 86.

The VTAM working set decreased from 584 pages in VM/ESA 1.0 370 Feature to 194 pages in VM/ESA 1.1. This apparent change results from the method of accounting for the GCS and VTAM shared pages. In VM/ESA 1.0 370 Feature, these pages are counted as part of the working set for the VTAM virtual machine. In VM/ESA 1.1, these pages are counted as part of the shared pages (SHRPGS), which increased from 379 to 1080 between these two measurements.

9221 TUNING RELEASE RUN ID	N/A ESA 1.0 (370) H17F0241	YES ESA 1.1 H14F0241
Environment		
REAL STORAGE	**64M**	**48M**
EXP. STORAGE	**0M**	**16M**
USERS	**240**	**240**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
Response Time		
TRIV INT	0.190	0.184
NONTRIV INT	3.272	1.034
TOT INT	0.550	0.789
TOT INT ADJ	0.508	0.668
AVG FIRST (T)	0.428	0.300
AVG LAST (T)	0.787	0.690
Throughput		
AVG THINK (T)	28.11	28.05
ETR	7.71	7.11
ETR (T)	8.35	8.40
ETR RATIO	0.923	0.847
ITR (H)	10.98	10.15
ITR	10.15	8.59
EMUL ITR	17.32	14.07
ITRR (H)	1.000	0.925
ITRR	1.000	0.846
Proc. Usage		
PBT/CMD (H)	91.101	98.483
PBT/CMD	91.015	98.853
CP/CMD (H)	40.212	44.457
CP/CMD	37.682	38.112
EMUL/CMD (H)	50.889	54.025
EMUL/CMD	53.320	60.741
Processor Util.		
TOTAL (H)	76.08	82.69
TOTAL	76.01	83.00
UTIL/PROC (H)	76.08	82.69
UTIL/PROC	76.01	83.00
TVR(H)	1.79	1.82
TVR	1.71	1.63
Storage		
WKSET (V)	119	82
PGBLPGS	15006	9623
PGBLPGS/USER	62.5	40.1
FREEPGS	818	730
FREE UTIL	na	0.87
SHRPGS	379	1080
Paging		
READS/SEC	38	58
WRITES/SEC	20	47
PAGE/CMD	6.922	12.505
XSTOR IN/SEC	na	0
XSTOR OUT/SEC	na	0
XSTOR/CMD	na	0.000
FAST CLR/CMD	na	5.717

9221 TUNING RELEASE RUN ID	N/A ESA 1.0 (370) H17F0241	YES ESA 1.1 H14F0241
Environment		
REAL STORAGE	**64M**	**48M**
EXP. STORAGE	**0M**	**16M**
USERS	**240**	**240**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
I/O		
VIO RATE	120	56
VIO/CMD	14.345	6.670
MDSK/CMD	3.951	na
MDC READS	na	47
MDC WRITES	na	16
MDC MODS	na	11
MDC HIT RATIO	na	0.89
PRIVOPs		
PRIVOP/CMD (R)	17.789	22.498
DIAG/CMD (R)	14.958	23.772
DIAG 08/CMD	na	0.715
DIAG 10/CMD	na	0.000
DIAG 58/CMD	na	1.191
DIAG 98/CMD	na	2.620
DIAG A4/CMD	na	2.501
DIAG A8/CMD	na	1.667
DIAG 214/CMD	na	10.838
SIE/CMD	na	69.554
SIE INTCPT/CMD	na	53.557
FREE TOTL/CMD	na	111.239
VTAM Machines		
WKSET (V)	583.5	194
TOT CPU/CMD (V)	13.5840	18.9415
CP CPU/CMD (V)	5.6684	8.5119
VIRT CPU/CMD (V)	7.9156	10.4296
DIAG 98/CMD (V)	2.603	2.710
SFS Servers		
WKSET (V)	675.75	401
TOT CPU/CMD (V)	8.8883	12.7847
CP CPU/CMD (V)	4.0920	6.8970
VIRT CPU/CMD (V)	4.7963	5.8877
FP REQ/CMD(Q)	1.329	1.340
IO/CMD (Q)	2.147	2.105
IO TIME/CMD (Q)	0.076	0.036
SFS TIME/CMD (Q)	0.114	0.051

Note: T=TPNS, V=VMPRF (ESA) or VMMAP (370), H=Hardware Monitor, R=RTM, Q=Query Filepool Counters, Unmarked=RTM (ESA) or VMMAP (370)

Table 13. 9221-170 / 35% SFS Comparing VM/ESA 1.1 to VM/ESA 1.0 370 Feature

Minidisk to Shared File System

The measurements in this section compare the performance of the CMS minidisk file system (EDF) to the Shared File System (SFS) to demonstrate the effects of migrating files from minidisk to SFS. For these measurements all end user data (i.e. all of the data accessed as Read/Write) was moved from minidisks to SFS.

9021-720 / Equal CPU Utilization

This section compares the performance of minidisk and SFS at similar processor utilization for a 9021-720.

1) WORKLOAD: FS7B35R and FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-720
- STORAGE:
   - RSTOR:   512M
   - XSTOR:   2G

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              PSPT01     3380-D
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
Minidisk:						
3380-A	15 - 3880-03	20	8	12	20	0
3380-D	3 - 3880-03	0	0	0	20	0
SFS:						
3380-A	10 - 3880-03	20	8	12	0	0
3380-K	4 - 3990-02	0	0	0	0	16

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

	CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
Minidisk:	3745-410	3	44	56Kb
SFS:	3745-410	3	22	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:    100
- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
Minidisk and SFS:				
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
SFS:				
SERVE2 (R/W)	SFS	32M/XA	1500	QUICKDSP ON
SERVE4 (R/W)	SFS	32M/XA	1500	QUICKDSP ON
SERVE7 (R/W)	SFS	32M/XA	1500	QUICKDSP ON
SERVE8 (R/W)	SFS	32M/XA	1500	QUICKDSP ON
CRRSERVA	CRR	32M/XA	100	NONE

4) MEASUREMENT DISCUSSION

Measurements were obtained for VM/ESA 1.0 and VM/ESA 1.1. Below are the major run characteristics for this comparison:

Y63R5866 VM/ESA 1.0 minidisk, with the number of users selected to obtain an approximate processor utilization of 90%.

Y63F4809 VM/ESA 1.0 35% SFS, with the number of users selected to obtain an approximate processor utilization of 90%.

Y64R5865 VM/ESA 1.1 minidisk, with the number of users set to the same number used for the VM/ESA 1.0 minidisk measurement.

Y64F480X VM/ESA 1.1 35% SFS, with the number of users set to the same number used for the VM/ESA 1.0 35% SFS measurement.

The VM/ESA 1.1 measurements showed that migrating all user files from minidisk to SFS while keeping the CPU utilization at approximately 90% decreased internal throughput (ITR (H)) by 14.7%. For VM/ESA 1.0, the internal throughput decrease was 15.3%.

Internal response time (TOT INT ADJ) was somewhat better for minidisk. For VM/ESA 1.0 the difference was 0.07 seconds, and for VM/ESA 1.1 the difference was 0.02 seconds. The VM/ESA 1.0 and VM/ESA 1.1 response times were influenced by the fact that the processor utilizations actually achieved were not quite the same and they departed from the 90% target. For example, if the VM/ESA 1.1 utilizations (87% for minidisk, 84% for SFS) were closer to 90%, the difference between minidisk and SFS would have been greater that 0.02 seconds.

The external response times (AVG LAST (T)) were longer for minidisk. This was due to the greater VTAM and network contention caused by the 5860 minidisk users (the SFS run only had 4800 users).

For VM/ESA 1.1 measurements, the processor busy time per command (PBT/CMD (H)) increased by 17.2% as the degree of SFS usage increased. The additional processor busy time per command was evenly split between CP and emulation. For the VM/ESA 1.0 measurements, the processor busy time per command (PBT/CMD (H)) increased by 18.1% as the degree of SFS usage increased. This demonstrated a narrowing of the minidisk-SFS processor usage gap between VM/ESA 1.0 and VM/ESA 1.1. Minidisk did improve between VM/ESA 1.0 and VM/ESA 1.1, but SFS improved even more due to the IUCV improvement which was more beneficial to SFS. For more information on the IUCV improvement see “IUCV Improvements” on page 10.

The measurements show that migrating data to SFS requires more system resources. The degree of increase is proportional to the amount of file I/O activity transferred from minidisk to SFS.

FILE SYSTEM RELEASE RUN ID	MINIDISK ESA 1.0 Y63R5866	35% SFS ESA 1.0 Y63F4809	MINIDISK ESA 1.1 Y64R5865	35% SFS ESA 1.1 Y64F480X
Environment				
REAL STORAGE	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**
USERS	**5860**	**4800**	**5860**	**4800**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**
Response Time				
TRIV INT	0.057	0.052	0.056	0.049
NONTRIV INT	0.328	0.437	0.303	0.341
TOT INT	0.192	0.262	0.193	0.220
TOT INT ADJ	0.251	0.317	0.224	0.245
AVG FIRST (T)	0.633	0.495	0.397	0.330
AVG LAST (T)	0.817	0.745	0.527	0.507
Throughput				
AVG THINK (T)	26.00	25.56	25.19	25.51
ETR	266.53	203.80	238.56	189.27
ETR (T)	204.06	168.70	205.97	169.82
ETR RATIO	1.306	1.208	1.158	1.115
ITR (H)	223.28	189.12	236.35	201.64
ITR	48.61	38.12	45.71	37.49
EMUL ITR	78.49	64.38	69.86	59.13
ITRR (H)	1.000	0.847	1.059	0.903
ITRR	1.000	0.784	0.940	0.771
Proc. Usage				
PBT/CMD (H)	26.872	31.726	25.386	29.755
PBT/CMD	26.855	31.713	25.344	29.737
CP/CMD (H)	10.640	13.365	9.268	11.427
CP/CMD	10.193	12.922	8.739	10.894
EMUL/CMD (H)	16.229	18.355	16.116	18.322
EMUL/CMD	16.662	18.791	16.604	18.843
Processor Util.				
TOTAL (H)	548.36	535.21	522.88	505.32
TOTAL	548.00	535.00	522.00	505.00
UTIL/PROC (H)	91.39	89.20	87.15	84.22
UTIL/PROC	91.33	89.17	87.00	84.17
TVR(H)	1.66	1.73	1.58	1.62
TVR	1.61	1.69	1.53	1.58
Storage				
WKSET (V)	54	61	58	64
PGBLPGS	105K	108K	104K	108K
PGBLPGS/USER	18.3	23.0	18.2	23.0
FREEPGS	13897	11511	14349	12217
FREE UTIL	0.96	0.96	0.96	0.96
SHRPGS	1195	1370	1186	1322
Paging				
READS/SEC	266	246	332	298
WRITES/SEC	161	154	186	161
PAGE/CMD	2.093	2.371	2.515	2.703
XSTOR IN/SEC	1053	1004	1053	1086
XSTOR OUT/SEC	1267	1207	1298	1298
XSTOR/CMD	11.369	13.106	11.414	14.038
FAST CLR/CMD	8.189	7.789	5.651	5.618

FILE SYSTEM RELEASE RUN ID	MINIDISK ESA 1.0 Y63R5866	35% SFS ESA 1.0 Y63F4809	MINIDISK ESA 1.1 Y64R5865	35% SFS ESA 1.1 Y64F480X
Environment				
REAL STORAGE	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**
USERS	**5860**	**4800**	**5860**	**4800**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**
I/O				
VIO RATE	1702	1151	1751	1178
VIO/CMD	8.341	6.823	8.501	6.937
MDC READS	1183	1014	1223	1044
MDC WRITES	536	304	550	309
MDC MODS	452	239	463	241
MDC HIT RATIO	0.93	0.94	0.93	0.93
PRIVOPs				
PRIVOP/CMD	19.015	27.636	20.075	28.826
DIAG/CMD	17.098	15.111	23.603	21.217
DIAG 08/CMD	0.720	0.741	0.719	0.730
DIAG 10/CMD	5.709	5.305	0.015	0.012
DIAG 58/CMD	1.225	1.239	1.219	1.242
DIAG 98/CMD	0.309	0.314	0.291	0.312
DIAG A4/CMD	3.945	2.608	4.083	2.726
DIAG A8/CMD	1.872	1.689	1.893	1.684
DIAG 214/CMD	na	na	12.371	11.547
SIE/CMD	65.235	78.910	54.688	72.357
SIE INTCPT/CMD	42.403	54.448	35.547	50.650
FREE TOTL/CMD	100.362	133.540	84.517	96.476
VTAM Machines				
WKSET (V)	1955	1529	1791	1402
TOT CPU/CMD (V)	4.3505	4.4557	3.8211	3.9843
CP CPU/CMD (V)	2.4394	2.5555	1.8823	1.9921
VIRT CPU/CMD (V)	1.9112	1.9002	1.9388	1.9921
DIAG 98/CMD (V)	0.313	0.319	0.295	0.316
SFS Servers				
WKSET (V)	na	1101	na	997
TOT CPU/CMD (V)	na	4.3404	na	3.7171
CP CPU/CMD (V)	na	2.4271	na	1.7960
VIRT CPU/CMD (V)	na	1.9133	na	1.9211
FP REQ/CMD(Q)	na	1.337	na	1.338
IO/CMD (Q)	na	2.002	na	1.982
IO TIME/CMD (Q)	na	0.041	na	0.036
SFS TIME/CMD (Q)	na	0.127	na	0.088
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM				

Table 14. Minidisk to SFS: 9021-720 / Equal CPU Utilization

9021-720 / Equal Number of Users

This section compares the performance of minidisk and SFS with an equal number of users.

1) WORKLOAD: FS7B35R and FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-720
- STORAGE:
   - RSTOR:   512M
   - XSTOR:   2G

- DASD:
             PACK NAME  TYPE
   - SYSTEM: PSYS02     3380-A
             PSPT01     3380-D
             WKLD01     3380-D
             WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
Minidisk:						
3380-A	15 - 3880-03	20	8	12	20	0
3380-D	3 - 3880-03	0	0	0	20	0
SFS:						
3380-A	10 - 3880-03	20	8	12	0	0
3380-K	4 - 3990-02	0	0	0	0	16

(NUMBER OF PACKS: PAGE, SPOOL, TDISK, USER, SERVER)

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

	CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
Minidisk:	3745-410	3	44	56Kb
SFS:	3745-410	3	22	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:    100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
Minidisk and SFS:				
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
SFS:				
SERVE2 (R/W)	SFS	32M/XA	1500	QUICKDSP ON
SERVE4 (R/W)	SFS	32M/XA	1500	QUICKDSP ON
SERVE7 (R/W)	SFS	32M/XA	1500	QUICKDSP ON
SERVE8 (R/W)	SFS	32M/XA	1500	QUICKDSP ON
CRRSERVA	CRR	32M/XA	100	NONE

4) MEASUREMENT DISCUSSION

This section summarizes the results of VM/ESA 1.1 measurements that compare minidisk and SFS with an equal number of users.

Real storage usage, response times, processor usage, and virtual file I/Os were higher for SFS than minidisk. The real storage and response time increases were larger than the corresponding increases for the equal utilization comparison (see section "9021-720 / Equal CPU Utilization" on page 94). The processor usage and virtual file I/O increases for the equal users measurements were similar to the corresponding increases for the equal utilization measurements.

A good measure of contention for real storage is the sum of PAGE/CMD and XSTOR/CMD. For minidisk the sum was 11.8 and for SFS the sum was 16.7. This was an increase of 42%. The increased real storage contention is due to:

- commands which use SFS require/reference more pages
- the servers also require real storage.

The corresponding increase for the VM/ESA 1.1 equal utilization measurements was only 20%. It was only 20% because the VM/ESA 1.1 equal utilization minidisk measurement had more users (5860) vying for real storage than this equal user minidisk measurement (4800 users).

Processor usage increased by 17.1% for SFS in these equal user measurements and by 17.2% in the equal utilization measurements.

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 virtual file I/Os per command for minidisk was 6.00 and for SFS was 6.39. This was an increase of 6.6%. The corresponding increase for the VM/ESA 1.1 equal utilization measurements was 7.0%.

Because of the increased processor usage and storage contention, external response time (AVG LAST (T)) increased by 0.17 seconds (49%). By contrast, the corresponding VM/ESA 1.1 equal utilization measurements showed a small decrease.

For both the equal utilization and equal users measurements, the minidisk cache (MDC) is equally effective for minidisk 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 minidisk 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 I/O, 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. Minidisk doesn't use block I/O, so this improvement does not apply.

2. SFS buffers are always 4K aligned (the case that MDC handles most efficiently). With minidisk, data may go directly to a user buffer which is not necessarily 4K aligned.

FILE SYSTEM RELEASE RUN ID	MINIDISK ESA 1.1 Y64R4801	35% SFS ESA 1.1 Y64F480X
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**4800**	**4800**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
Response Time		
TRIV INT	0.044	0.049
NONTRIV INT	0.227	0.341
TOT INT	0.155	0.220
TOT INT ADJ	0.164	0.245
AVG FIRST (T)	0.230	0.330
AVG LAST (T)	0.340	0.507
Throughput		
AVG THINK (T)	25.88	25.51
ETR	180.60	189.27
ETR (T)	170.35	169.82
ETR RATIO	1.060	1.115
ITR (H)	236.15	201.64
ITR	41.78	37.49
EMUL ITR	63.93	59.13
ITRR (H)	1.000	0.854
ITRR	1.000	0.897
Proc. Usage		
PBT/CMD (H)	25.407	29.755
PBT/CMD	25.360	29.737
CP/CMD (H)	9.248	11.427
CP/CMD	8.806	10.894
EMUL/CMD (H)	16.156	18.322
EMUL/CMD	16.554	18.843
Processor Util.		
TOTAL (H)	432.81	505.32
TOTAL	432.00	505.00
UTIL/PROC (H)	72.14	84.22
UTIL/PROC	72.00	84.17
TVR(H)	1.57	1.62
TVR	1.53	1.58
Storage		
WKSET (V)	60	64
PGBLPGS	109K	108K
PGBLPGS/USER	23.3	23.0
FREEPGS	11767	12217
FREE UTIL	0.96	0.96
SHRPGS	1111	1322
Paging		
READS/SEC	223	298
WRITES/SEC	74	161
PAGE/CMD	1.744	2.703
XSTOR IN/SEC	801	1086
XSTOR OUT/SEC	910	1298
XSTOR/CMD	10.044	14.038
FAST CLR/CMD	5.630	5.618

FILE SYSTEM RELEASE RUN ID	MINIDISK ESA 1.1 Y64R4801	35% SFS ESA 1.1 Y64F480X
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**4800**	**4800**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
I/O		
VIO RATE	1454	1178
VIO/CMD	8.536	6.937
MDC READS	1006	1044
MDC WRITES	451	309
MDC MODS	384	241
MDC HIT RATIO	0.93	0.93
PRIVOPs		
PRIVOP/CMD	21.459	28.826
DIAG/CMD	23.935	21.217
DIAG 08/CMD	0.751	0.730
DIAG 10/CMD	0.012	0.012
DIAG 58/CMD	1.239	1.242
DIAG 98/CMD	0.305	0.312
DIAG A4/CMD	4.068	2.726
DIAG A8/CMD	1.931	1.684
DIAG 214/CMD	12.422	11.547
SIE/CMD	56.913	72.357
SIE INTCPT/CMD	38.701	50.650
FREE TOTL/CMD	84.158	96.476
VTAM Machines		
WKSET (V)	1432	1402
TOT CPU/CMD (V)	4.0385	3.9843
CP CPU/CMD (V)	2.0405	1.9921
VIRT CPU/CMD (V)	1.9981	1.9921
DIAG 98/CMD (V)	0.307	0.316
SFS Servers		
WKSET (V)	na	997
TOT CPU/CMD (V)	na	3.7171
CP CPU/CMD (V)	na	1.7960
VIRT CPU/CMD (V)	na	1.9211
FP REQ/CMD(Q)	na	1.338
IO/CMD (Q)	na	1.982
IO TIME/CMD (Q)	na	0.036
SFS TIME/CMD (Q)	na	0.088
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM		

Table 15. Minidisk to SFS: 9021-720 / Equal Number of Users

9121-320 / Equal CPU Utilization

This section compares the performance of minidisk and SFS at similar processor utilization for a 9121-320.

1) WORKLOAD: FS7B0R and FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9121-320
- STORAGE:
   - RSTOR:       192M
   - XSTOR:        64M     All reserved for MDC

- DASD:
              PACK NAME   TYPE
   - SYSTEM:  PSYS02      3380-A
              PSPT01      3380-A
              WKLD01      3380-A
              WKLD02      3380-A
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-A	11 - 3880-2	16	4	8	0	10
3380-A	2 - 3880-J23	0	0	0	0	4

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	1	44	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:   100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
Minidisk and SFS:				
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
SFS:				
RWSERV1	SFS	32M/XA	1500	QUICKDSP ON
CRRSERV1	CRR	17M/XA	100	NONE

4) MEASUREMENT DISCUSSION

Measurements were obtained for VM/ESA 1.1. No expanded storage was used for paging (for an explanation of this, see "Using XSTOR on a 9121" on page 43).

Results were similar to "9021-720 / Equal CPU Utilization" on page 94, except for the internal response time (TOT INT ADJ) increase between minidisk and SFS. For the 9021-720 measurements internal response times were similar, but for these 9121-320 measurements there was an increase of 0.14 seconds (43%). This difference is because the 9121-320 runs had more similar processor utilizations and because the 9121-320 runs had no XSTOR for paging. As discussed in "9021-720 / Equal CPU Utilization" on page 94, if the processor utilizations had been closer, the response time difference would have been greater.

With no XSTOR for paging (XSTOR/CMD) in the 9121-320 measurements, there was more DASD paging (PAGE/CMD) than in the 9021-720 measurements. This influences the SFS 9121-320 internal response time (TOT INT ADJ) because SFS in general has greater storage requirements than minidisk and because of the interaction of the SFS server checkpoint processing and paging to DASD. During checkpoint processing, server requests are queued as they continue to arrive. This queueing causes additional page references. The resulting DASD paging elongates the checkpoint processing and the increased checkpoint serialization causes the average response time to get longer.

FILE SYSTEM RELEASE RUN ID	MINIDISK ESA 1.1 L14R0910	35% SFS ESA 1.1 L14F0770
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**910**	**770**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**1**	**1**
Response Time		
TRIV INT	0.068	0.069
NONTRIV INT	0.453	0.656
TOT INT	0.319	0.452
TOT INT ADJ	0.329	0.469
AVG FIRST (T)	0.280	0.330
AVG LAST (T)	0.490	0.640
Throughput		
AVG THINK (T)	25.82	25.37
ETR	33.39	28.59
ETR (T)	32.41	27.58
ETR RATIO	1.030	1.037
ITR (H)	36.94	31.56
ITR	38.09	32.73
EMUL ITR	58.55	51.60
ITRR (H)	1.000	0.854
ITRR	1.000	0.859
Proc. Usage		
PBT/CMD (H)	27.070	31.688
PBT/CMD	27.148	31.546
CP/CMD (H)	9.784	11.900
CP/CMD	9.564	11.603
EMUL/CMD (H)	17.279	19.777
EMUL/CMD	17.585	19.943
Processor Util.		
TOTAL (H)	87.75	87.39
TOTAL	88.00	87.00
UTIL/PROC (H)	87.75	87.39
UTIL/PROC	88.00	87.00
TVR(H)	1.57	1.60
TVR	1.54	1.58
Storage		
WKSET (V)	81	88
PGBLPGS	43242	43749
PGBLPGS/USER	47.5	56.8
FREEPGS	2373	2117
FREE UTIL	0.93	0.92
SHRPGS	784	1038
Paging		
READS/SEC	252	224
WRITES/SEC	159	147
PAGE/CMD	12.679	13.452
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.399	5.584

FILE SYSTEM RELEASE RUN ID	MINIDISK ESA 1.1 L14R0910	35% SFS ESA 1.1 L14F0770
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**910**	**770**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**1**	**1**
I/O		
VIO RATE	279	193
VIO/CMD	8.607	6.998
MDC READS	187	162
MDC WRITES	88	52
MDC MODS	72	37
MDC HIT RATIO	0.92	0.91
PRIVOPs		
PRIVOP/CMD	24.160	32.175
DIAG/CMD	25.339	22.826
DIAG 08/CMD	0.771	0.725
DIAG 10/CMD	0.000	0.000
DIAG 58/CMD	1.234	1.233
DIAG 98/CMD	0.432	0.508
DIAG A4/CMD	4.041	2.719
DIAG A8/CMD	2.005	1.740
DIAG 214/CMD	12.340	11.567
SIE/CMD	61.485	73.680
SIE INTCPT/CMD	44.884	55.997
FREE TOTL/CMD	86.720	98.953
VTAM Machines		
WKSET (V)	1055	542
TOT CPU/CMD (V)	4.8675	4.9152
CP CPU/CMD (V)	2.4337	2.4979
VIRT CPU/CMD (V)	2.4337	2.4173
DIAG 98/CMD (V)	0.461	0.529
SFS Servers		
WKSET (V)	na	1463
TOT CPU/CMD (V)	na	3.8476
CP CPU/CMD (V)	na	1.8331
VIRT CPU/CMD (V)	na	2.0144
FP REQ/CMD(Q)	na	1.357
IO/CMD (Q)	na	1.888
IO TIME/CMD (Q)	na	0.031
SFS TIME/CMD (Q)	na	0.096
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM		

Table 16. Minidisk to SFS: 9121-320 / Equal CPU Utilization

9221-170 / Equal CPU Utilization

This section compares the performance of minidisk and SFS at similar processor utilization on a 9221-170.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9221-170
- STORAGE:
   CP/370 - RSTOR:   64M
          - XSTOR:   0M
   ESA
          - RSTOR:   48M
          - XSTOR:   16M (all reserved for MDC)

- DASD:
             PACK NAME   TYPE
   - SYSTEM: PRF05       3380
     CP/370  PRF01       3380
             PRFRES      3380

   - SYSTEM: H3AP01      3380
     ESA     H3SRV       3380
             H3RES       3380
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-A	3 - 3380-03	2	2	2	4	5
3380-D	1 - 3380-03	1	0	1	1	1

(NUMBER OF PACKS: PAGE, SPOOL, TDISK, USER, SERVER)

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINESPEED
3088-02	1	4.5M

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR:  BACTRIAN
- CMS BLOCKSIZE:     4K
- USER VM SIZE:      2M
- USER CMS MODE:     370
- USER RELSHARE:      64  CP/370
- USER RELSHARE:     100  ESA
- SERVER MACHINES:
```

	SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
CP/370					
	VTAM	VTAM/VSCS	16M/370		PRIORITY 1, QDROP OFF USERS, FAVOR 100, FAVOR
	RWSERVE1 (R/W)	SFS	16M/370		PRIORITY 1, QDROP OFF USERS
	CRRRECOV	CRR	16M/370		NONE
ESA					
	VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON
	RWSERVE1 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
	CRRRECOV	CRR	32M/XC	100	NONE

4) MEASUREMENT DISCUSSION

This section compares the performance of minidisk versus 35% SFS at similar processor utilization on a 9221-170 for VM/ESA 1.0 370 Feature and VM/ESA 1.1. The VM/ESA 1.1 measurements have been tuned. See section “Recommended 9221 Tuning” on page 223 for more details on tuning for VM/ESA 1.1.

Comparing the VM/ESA 1.0 370 Feature minidisk measurement (H17R0281) to the VM/ESA 1.0 370 Feature 35% SFS measurement (H17F0241), internal throughput (ITR(H)) decreased by 12.8%. External response time (AVG LAST (T)) improved by 1.6%. The processor busy time per command (PBT/CMD) increased by 14.7%.

Comparing VM/ESA 1.1 minidisk measurement (H14R0287) to the VM/ESA 1.1 35% SFS measurement (H14F0241), internal throughput decreased by 14.6%. External response time increased by 25.5%. The processor busy time per command increased by 17.2%. The internal throughput decrease on VM/ESA 1.1 (14.6%) was larger than the internal throughput decrease on VM/ESA 1.0 370 Feature (12.8%) because SFS makes extensive use of APPC/VM and block I/O, which have longer pathlengths on VM/ESA 1.1.

With no XSTOR for paging (XSTOR/CMD), there was more DASD paging (PAGE/CMD) as the case with the 9121-320 measurements (see section “9121-320 / Equal CPU Utilization” on page 103). This influenced the SFS 9221-170 internal response time (TOT INT ADJ) because SFS in general has greater storage requirements than minidisk and because of the interaction of the SFS server checkpoint processing and paging to DASD. During checkpoint processing, server requests are queued as they continue to arrive. This queueing caused additional page references. The resulting DASD paging elongated the checkpoint processing and the increased checkpoint serialization caused the average response time to get longer.

FILE SYSTEM **RELEASE** **RUN ID**	**MINIDISK** **ESA 1.0 (370)** **H17R0281**	**35% SFS** **ESA 1.0 (370)** **H17F0241**	**MINDISK** **ESA 1.1** **H14R0287**	**35% SFS** **ESA 1.1** **H14F0241**
Environment				
REAL STORAGE	**64M**	**64M**	**48M**	**48M**
EXP. STORAGE	**0M**	**0M**	**16M**	**16M**
USERS	**280**	**240**	**280**	**240**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**	**0**
PROCESSORS	**1**	**1**	**1**	**1**
Response Time				
TRIV INT	0.180	0.190	0.172	0.184
NONTRIV INT	3.140	3.272	0.810	1.034
TOT INT	0.500	0.550	0.624	0.789
TOT INT ADJ	0.461	0.508	0.529	0.668
AVG FIRST (T)	0.460	0.428	0.270	0.300
AVG LAST (T)	0.800	0.787	0.550	0.690
Throughput				
AVG THINK (T)	28.57	28.11	28.39	28.05
ETR	8.79	7.71	8.23	7.11
ETR (T)	9.53	8.35	9.71	8.40
ETR RATIO	0.922	0.923	0.848	0.847
ITR (H)	12.59	10.98	11.90	10.15
ITR	11.62	10.15	10.10	8.59
EMUL ITR	19.71	17.32	15.99	14.07
ITRR (H)	1.000	0.872	0.945	0.806
ITRR	1.000	0.873	0.869	0.739
Proc. Usage				
PBT/CMD (H)	79.416	91.101	84.060	98.483
PBT/CMD	79.297	91.015	84.486	98.853
CP/CMD (H)	34.756	40.212	36.380	44.457
CP/CMD	32.533	37.682	30.910	38.112
EMUL/CMD (H)	44.659	50.889	47.680	54.025
EMUL/CMD	46.765	53.320	53.577	60.741
Processor Util.				
TOTAL (H)	75.72	76.08	81.59	82.69
TOTAL	75.61	76.01	82.00	83.00
UTIL/PROC (H)	75.72	76.08	81.59	82.69
UTIL/PROC	75.61	76.01	82.00	83.00
TVR(H)	1.78	1.79	1.76	1.82
TVR	1.70	1.71	1.58	1.63
Storage				
WKSET (V)	90	119	77	82
PGBLPGS	15006	15006	9520	9623
PGBLPGS/USER	53.6	62.5	34.0	40.1
FREEPGS	864	818	804	730
FREE UTIL	na	na	0.88	0.87
SHRPGS	286	379	900	1080
Paging				
READS/SEC	43	38	62	58
WRITES/SEC	23	20	49	47
PAGE/CMD	6.857	6.922	11.437	12.505
XSTOR IN/SEC	na	na	0	0
XSTOR OUT/SEC	na	na	0	0
XSTOR/CMD	na	na	0.000	0.000
FAST CLR/CMD	na	na	5.461	5.717

FILE SYSTEM RELEASE RUN ID	MINIDISK ESA 1.0 (370) H17R0281	35% SFS ESA 1.0 (370) H17F0241	MINDISK ESA 1.1 H14R0287	35% SFS ESA 1.1 H14F0241
Environment				
REAL STORAGE	**64M**	**64M**	**48M**	**48M**
EXP. STORAGE	**0M**	**0M**	**16M**	**16M**
USERS	**280**	**240**	**280**	**240**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**	**0**
PROCESSORS	**1**	**1**	**1**	**1**
I/O				
VIO RATE	144	120	82	56
VIO/CMD	15.145	14.345	8.449	6.670
MDSK/CMD	5.558	3.951	na	na
MDC READS	na	na	53	47
MDC WRITES	na	na	26	16
MDC MODS	na	na	21	11
MDC HIT RATIO	na	na	0.91	0.89
PRIVOPs				
PRIVOP/CMD (R)	9.077	17.789	14.426	22.498
DIAG/CMD (R)	16.375	14.958	26.002	23.772
DIAG 08/CMD	na	na	0.618	0.715
DIAG 10/CMD	na	na	0.000	0.000
DIAG 58/CMD	na	na	1.236	1.191
DIAG 98/CMD	na	na	2.473	2.620
DIAG A4/CMD	na	na	3.812	2.501
DIAG A8/CMD	na	na	2.061	1.667
DIAG 214/CMD	na	na	11.334	10.838
SIE/CMD	na	na	56.977	69.554
SIE INTCPT/CMD	na	na	42.163	53.557
FREE TOTL/CMD	na	na	97.571	111.239
VTAM Machines				
WKSET (V)	529	583.5	211	194
TOT CPU/CMD (V)	11.8812	13.5840	18.3362	18.9415
CP CPU/CMD (V)	4.7818	5.6684	8.1785	8.5119
VIRT CPU/CMD (V)	7.0994	7.9156	10.1577	10.4296
DIAG 98/CMD (V)	1.834	2.603	2.519	2.710
SFS Servers				
WKSET (V)	na	675.75	na	401
TOT CPU/CMD (V)	na	8.8883	na	12.7847
CP CPU/CMD (V)	na	4.0920	na	6.8970
VIRT CPU/CMD (V)	na	4.7963	na	5.8877
FP REQ/CMD(Q)	na	1.329	na	1.340
IO/CMD (Q)	na	2.147	na	2.105
IO TIME/CMD (Q)	na	0.076	na	0.036
SFS TIME/CMD (Q)	na	0.114	na	0.051

Note: T=TPNS, V=VMPRF (ESA) or VMMAP (370), H=Hardware Monitor, R=RTM, Q=Query Filepool Counters, Unmarked=RTM (ESA) or VMMAP (370)

Table 17. 9221-170 / Equal CPU Utilization Comparing VM/ESA 1.0 370 Feature and VM/ESA 1.1

Virtual Machine Storage Considerations

The measurements in this section deal with virtual machine storage considerations. Changes in the storage size of virtual machines and the placement of saved segments can impact system performance. See “Virtual Machine Storage Considerations” on page 27 for additional information and associated concepts.

3090-300J / Virtual Machine Size

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  3090-300J
- STORAGE:
   - RSTOR:        256M
   - XSTOR:          1G

- DASD:
              PACK NAME   TYPE
   - SYSTEM:  RESPAK      3380-A
              SRVPAK      3380-A
              ESAP01      3380-A
              ESAOV1      3380-A
              ESAOV2      3380-A
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3390-A	2 - 3990-3	5	5	5	0	10
3380-A	1 - 3880-2	0	0	0	5	0

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CTCA	NUMBER	CHANNEL SPEED
3088	1	4.5M

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR:  BACTRIAN
- CMS BLOCKSIZE:     4K
- USER VM SIZE:      2M or 4M (See Table)
- USER CMS MODE:     XA
- USER RELSHARE:     100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON
RWSERV1	SFS	32M/XA	1500	QUICKDSP ON
RWSERV2	SFS	32M/XA	1500	QUICKDSP ON
CRRSERV1	CRR	17M/XA	100	

4) MEASUREMENT DISCUSSION

The following table shows measurements where the virtual machine storage size was varied from 2M to 4M. This was expected to cause a small increase in paging due to additional references for CMS Storage Management control blocks.

The results showed system performance to be equivalent between the two measurements. The increases in paging (PAGE/CMD) and user working set size (WKSET) were negligible. This illustrates that virtual machine storage sizes can be increased without impacting system performance. However, there are scenarios where this is not true (see "Virtual Machine Storage Considerations" on page 27 for details). These include the following:

- Use of virtual machines with storage sizes greater than 32M
- Applications or products that behave differently based on the amount of virtual storage available.

VM SIZE RELEASE RUN ID	2M ESA 1.1 Y34F2644	4M ESA 1.1 Y34F2645
Environment		
REAL STORAGE	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**
USERS	**2640**	**2640**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**3**	**3**
Response Time		
TRIV INT	0.108	0.108
NONTRIV INT	0.521	0.530
TOT INT	0.389	0.395
TOT INT ADJ	0.338	0.344
AVG FIRST (T)	0.210	0.213
AVG LAST (T)	0.387	0.393
Throughput		
AVG THINK (T)	25.46	25.28
ETR	82.71	83.28
ETR (T)	95.27	95.67
ETR RATIO	0.868	0.870
ITR (H)	108.78	108.73
ITR	31.54	31.65
EMUL ITR	47.63	47.93
ITRR (H)	1.000	1.000
ITRR	1.000	1.004
Proc. Usage		
PBT/CMD (H)	27.580	27.590
PBT/CMD	27.501	27.490
CP/CMD (H)	9.817	9.872
CP/CMD	9.237	9.303
EMUL/CMD (H)	17.759	17.714
EMUL/CMD	18.264	18.188
Processor Util.		
TOTAL (H)	262.75	263.96
TOTAL	262.00	263.00
UTIL/PROC (H)	87.58	87.99
UTIL/PROC	87.33	87.67
TVR(H)	1.55	1.56
TVR	1.51	1.51
Storage		
WKSET (V)	64	65
PGBLPGS	49124	49236
PGBLPGS/USER	18.6	18.7
FREEPGS	6677	6674
FREE UTIL	0.96	0.96
SHRPGS	1196	1165
Paging		
READS/SEC	218	218
WRITES/SEC	77	81
PAGE/CMD	3.096	3.125
XSTOR IN/SEC	678	676
XSTOR OUT/SEC	778	786
XSTOR/CMD	15.283	15.282
FAST CLR/CMD	5.563	5.634

VM SIZE RELEASE RUN ID	2M ESA 1.1 Y34F2644	4M ESA 1.1 Y34F2645
Environment		
REAL STORAGE	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**
USERS	**2640**	**2640**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**3**	**3**
I/O		
VIO RATE	660	663
VIO/CMD	6.928	6.930
MDC READS	526	526
MDC WRITES	176	180
MDC MODS	111	115
MDC HIT RATIO	0.89	0.88
PRIVOPs		
PRIVOP/CMD	24.864	24.979
DIAG/CMD	22.320	22.356
DIAG 08/CMD	0.735	0.742
DIAG 10/CMD	0.010	0.010
DIAG 58/CMD	1.249	1.254
DIAG 98/CMD	0.756	0.753
DIAG A4/CMD	2.750	2.739
DIAG A8/CMD	1.627	1.662
DIAG 214/CMD	11.536	11.550
SIE/CMD	67.734	67.754
SIE INTCPT/CMD	45.382	45.395
FREE TOTL/CMD	95.371	95.600
VTAM Machines		
WKSET (V)	925	878
TOT CPU/CMD (V)	3.1023	3.1009
CP CPU/CMD (V)	1.4112	1.4111
VIRT CPU/CMD (V)	1.6911	1.6898
DIAG 98/CMD (V)	0.759	0.753
SFS Servers		
WKSET (V)	988	1024
TOT CPU/CMD (V)	3.6096	3.5307
CP CPU/CMD (V)	1.6211	1.6143
VIRT CPU/CMD (V)	1.9885	1.9163
FP REQ/CMD(Q)	1.317	1.312
IO/CMD (Q)	1.803	1.832
IO TIME/CMD (Q)	0.033	0.034
SFS TIME/CMD (Q)	0.095	0.099
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM		

Table 18. The Effect of Increasing the Virtual Machine Storage Size

3090-300J / Placement of Saved Segments

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

- PROCESSOR: 3090-300J
- STORAGE:
  - RSTOR: 256M
  - XSTOR: 1G

- DASD:

	PACK NAME	TYPE
- SYSTEM:	RESPAK	3380-A
	SRVPAK	3380-A
	ESAP01	3380-A
	ESAOV1	3380-A
	ESAOV2	3380-A

TYPE OF	NUMBER/TYPE OF	NUMBER OF PACKS				
DASD	CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3390-A	2 - 3990-3	5	5	5	0	10
3380-A	1 - 3880-2	0	0	0	5	0

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CTCA	NUMBER	CHANNEL SPEED
3088	1	4.5M

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2M
- USER CMS MODE: XA
- USER RELSHARE: 100

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON
RWSERV1	SFS	32M/XA	1500	QUICKDSP ON
RWSERV2	SFS	32M/XA	1500	QUICKDSP ON
CRRSERV1	CRR	17M/XA	100	

4) MEASUREMENT DISCUSSION

In the following three measurements, the location of the Callable Services Library (CSL) saved segment was moved from 7M to 30M to 35M. This segment is named VMLIB. The impact of moving the segment is the need for CP segment tables for the user virtual machines. The segment table fits inside the Virtual Machine Definition Block (VMDBK) when VMLIB is located below 32M. For addressability above 32M, a separate 4K segment table is required. Segment tables are not eligible for paging.

Comparison of the measurements with VMLIB at 7M and 30M showed no difference in system performance. In both cases, the segment table fit inside the VMDBK so real storage requirements remain the same. This illustrates that saved segments can usually be moved around below the 32M line without impacting system performance.

System performance remained equivalent when comparing the measurements with VMLIB at 7M and 35M. However, there was a significant system change. The number of pageable pages (PGBLPGS) is significantly reduced with VMLIB at 35M. This is caused by the need for a segment table separate from the VMDBK. If normalized per user, the result is one less pageable page per user. This maps directly with need for one additional non-pageable page per user for the segment table. In this configuration, the loss of one page is not significant and therefore did not impact system performance. See “Storage Constrained Runs” on page 34 for more information on where this change could impact system performance.

VMLIB ADDRESS RELEASE RUN ID	7M ESA 1.1 Y34F2644	30M ESA 1.1 Y34F2647	35M ESA 1.1 Y34F2648
Environment			
REAL STORAGE	**256M**	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**	**1024M**
USERS	**2640**	**2640**	**2640**
VTAMs	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**
PROCESSORS	**3**	**3**	**3**
Response Time			
TRIV INT	0.108	0.107	0.108
NONTRIV INT	0.521	0.529	0.529
TOT INT	0.389	0.394	0.395
TOT INT ADJ	0.338	0.343	0.344
AVG FIRST (T)	0.210	0.217	0.210
AVG LAST (T)	0.387	0.393	0.390
Throughput			
AVG THINK (T)	25.46	25.43	25.45
ETR	82.71	82.99	82.94
ETR (T)	95.27	95.46	95.20
ETR RATIO	0.868	0.869	0.871
ITR (H)	108.78	109.52	109.39
ITR	31.54	31.81	31.74
EMUL ITR	47.63	48.19	48.05
ITRR (H)	1.000	1.007	1.006
ITRR	1.000	1.009	1.007
Proc. Usage			
PBT/CMD (H)	27.580	27.392	27.425
PBT/CMD	27.501	27.342	27.417
CP/CMD (H)	9.817	9.791	9.769
CP/CMD	9.237	9.324	9.244
EMUL/CMD (H)	17.759	17.596	17.652
EMUL/CMD	18.264	18.019	18.173
Processor Util.			
TOTAL (H)	262.75	261.47	261.08
TOTAL	262.00	261.00	261.00
UTIL/PROC (H)	87.58	87.16	87.03
UTIL/PROC	87.33	87.00	87.00
TVR(H)	1.55	1.56	1.55
TVR	1.51	1.52	1.51
Storage			
WKSET (V)	64	65	64
PGBLPGS	49124	49329	46782
PGBLPGS/USER	18.6	18.7	17.7
FREEPGS	6677	6678	6689
FREE UTIL	0.96	0.96	0.96
SHRPGS	1196	1172	1125
Paging			
READS/SEC	218	226	217
WRITES/SEC	77	89	83
PAGE/CMD	3.096	3.300	3.151
XSTOR IN/SEC	678	667	676
XSTOR OUT/SEC	778	779	781
XSTOR/CMD	15.283	15.148	15.305
FAST CLR/CMD	5.563	5.573	5.536

VMLIB ADDRESS RELEASE RUN ID	7M ESA 1.1 Y34F2644	30M ESA 1.1 Y34F2647	35M ESA 1.1 Y34F2648
Environment			
REAL STORAGE	**256M**	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**	**1024M**
USERS	**2640**	**2640**	**2640**
VTAMs	**1**	**1**	**1**
VSCSs	**0**	**0**	**0**
PROCESSORS	**3**	**3**	**3**
I/O			
VIO RATE	660	660	662
VIO/CMD	6.928	6.914	6.954
MDC READS	526	530	524
MDC WRITES	176	178	177
MDC MODS	111	119	113
MDC HIT RATIO	0.89	0.89	0.89
PRIVOPs			
PRIVOP/CMD	24.864	24.916	24.901
DIAG/CMD	22.320	22.291	22.296
DIAG 08/CMD	0.735	0.733	0.735
DIAG 10/CMD	0.010	0.010	0.011
DIAG 58/CMD	1.249	1.247	1.261
DIAG 98/CMD	0.756	0.765	0.756
DIAG A4/CMD	2.750	2.734	2.731
DIAG A8/CMD	1.627	1.645	1.702
DIAG 214/CMD	11.536	11.534	11.502
SIE/CMD	67.734	67.486	67.597
SIE INTCPT/CMD	45.382	45.216	45.290
FREE TOTL/CMD	95.371	93.948	93.889
VTAM Machines			
WKSET (V)	925	766	834
TOT CPU/CMD (V)	3.1023	3.0962	3.1105
CP CPU/CMD (V)	1.4112	1.4084	1.4181
VIRT CPU/CMD (V)	1.6911	1.6878	1.6924
DIAG 98/CMD (V)	0.759	0.771	0.761
SFS Servers			
WKSET (V)	988	947	1061
TOT CPU/CMD (V)	3.6096	3.4454	3.5949
CP CPU/CMD (V)	1.6211	1.5888	1.5932
VIRT CPU/CMD (V)	1.9885	1.8566	2.0017
FP REQ/CMD(Q)	1.317	1.321	1.308
IO/CMD (Q)	1.803	1.836	1.824
IO TIME/CMD (Q)	0.033	0.033	0.034
SFS TIME/CMD (Q)	0.095	0.098	0.097
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM			

Table 19. The Effect of Moving the CSL Saved Segment

Software Mode Comparisons

9021-720 / 35% SFS

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-720
- STORAGE:
   - RSTOR:       512M
   - XSTOR:      2048M

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              PSPT01     3380-D
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	10 - 3880-3	20	8	12	0	0
3380-K	4 - 3990-2	0	0	0	0	16

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	22	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:             TPNS
- THINK TIME DISTR:   BACTRIAN
- CMS BLOCKSIZE:      4K
- USER VM SIZE:       2M
- USER CMS MODE:      370, XA, XC
- USER RELSHARE:      100
```

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
SERVE2	SFS	32M/XA	1500	QUICKDSP ON
SERVE4	SFS	32M/XA	1500	QUICKDSP ON
SERVE7	SFS	32M/XA	1500	QUICKDSP ON
SERVE8	SFS	32M/XA	1500	QUICKDSP ON
CRRSERVA	CRR	16M/XA	100	

4) MEASUREMENT DISCUSSION

These measurements were made to examine the performance effects of changing the user virtual machine modes. The key difference should be realized in the CPU resources consumed. In VM/ESA 1.0, when going from 370 mode to XA mode, the processor busy time per command (PBT/CMD (H)) increased by 3.0% for the measured environment. In VM/ESA 1.1 this increase was reduced to 1.2%. This was accomplished by reducing the number of instructions executed in XA mode unique paths, primarily in the SVC interrupt handler. When running XC mode in VM/ESA 1.1, there was an additional 0.9% of CPU time required to support this new environment.

A summary of virtual pathlength traces made in each of these environments for selected commands can be found in Appendix A. The number of Special Operations, which include those assisted privileged instructions required to support the various modes, decreased. VM/ESA 1.0 required an average of 116 (72%) more special operations to support XA mode, while VM/ESA 1.1 required an average of 57 (36%) more to support the XA mode environment.

USER MODE RELEASE RUN ID	370 ESA 1.0 Y63F480B	XA ESA 1.0 Y63F4809	370 ESA 1.1 Y64F480M	XA ESA 1.1 Y64F480X	XC ESA 1.1 Y64F480L
Environment					
REAL STORAGE	**512M**	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**	**2048M**
USERS	**4800**	**4800**	**4800**	**4800**	**4800**
VTAMs	**1**	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**	**6**
Response Time					
TRIV INT	0.054	0.052	0.052	0.049	0.049
NONTRIV INT	0.450	0.437	0.359	0.341	0.346
TOT INT	0.284	0.262	0.238	0.220	0.222
TOT INT ADJ	0.325	0.317	0.259	0.245	0.249
AVG FIRST (T)	0.447	0.495	0.333	0.330	0.357
AVG LAST (T)	0.675	0.745	0.497	0.507	0.525
Throughput					
AVG THINK (T)	25.57	25.56	25.66	25.51	25.61
ETR	193.42	203.80	185.18	189.27	190.35
ETR (T)	168.83	168.70	169.93	169.82	169.39
ETR RATIO	1.146	1.208	1.090	1.115	1.124
ITR (H)	194.83	189.12	204.15	201.64	199.92
ITR	37.15	38.12	37.13	37.49	37.48
EMUL ITR	63.40	64.38	59.33	59.13	58.88
ITRR (H)	1.000	0.971	1.048	1.035	1.026
ITRR	1.000	1.026	0.999	1.009	1.009
Proc. Usage					
PBT/CMD (H)	30.796	31.726	29.391	29.755	30.012
PBT/CMD	30.860	31.713	29.424	29.737	29.991
CP/CMD (H)	13.209	13.365	11.537	11.427	11.445
CP/CMD	12.794	12.922	11.005	10.894	10.922
EMUL/CMD (H)	17.581	18.355	17.847	18.322	18.561
EMUL/CMD	18.066	18.791	18.420	18.843	19.069
Processor Util.					
TOTAL (H)	519.92	535.21	499.42	505.32	508.36
TOTAL	521.00	535.00	500.00	505.00	508.00
UTIL/PROC (H)	86.65	89.20	83.24	84.22	84.73
UTIL/PROC	86.83	89.17	83.33	84.17	84.67
TVR(H)	1.75	1.73	1.65	1.62	1.62
TVR	1.71	1.69	1.60	1.58	1.57
Storage					
WKSET (V)	60	61	64	64	64
PGBLPGS	109K	108K	108K	108K	108K
PGBLPGS/USER	23.3	23.0	23.0	23.0	23.0
FREEPGS	11498	11511	12189	12217	12229
FREE UTIL	0.96	0.96	0.97	0.96	0.96
SHRPGS	1330	1370	1362	1322	1354
Paging					
READS/SEC	236	246	302	298	298
WRITES/SEC	141	154	173	161	163
PAGE/CMD	2.233	2.371	2.795	2.703	2.722
XSTOR IN/SEC	997	1004	1086	1086	1086
XSTOR OUT/SEC	1193	1207	1313	1298	1304
XSTOR/CMD	12.972	13.106	14.118	14.038	14.110
FAST CLR/CMD	7.570	7.789	5.467	5.618	5.632

USER MODE **RELEASE** **RUN ID**	**370** **ESA 1.0** **Y63F480B**	**XA** **ESA 1.0** **Y63F4809**	**370** **ESA 1.1** **Y64F480M**	**XA** **ESA 1.1** **Y64F480X**	**XC** **ESA 1.1** **Y64F480L**
Environment					
REAL STORAGE	**512M**	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**	**2048M**
USERS	**4800**	**4800**	**4800**	**4800**	**4800**
VTAMs	**1**	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**	**6**
I/O					
VIO RATE	1144	1151	1176	1178	1174
VIO/CMD	6.776	6.823	6.921	6.937	6.931
MDC READS	1012	1014	1052	1044	1041
MDC WRITES	308	304	312	309	310
MDC MODS	240	239	242	241	241
MDC HIT RATIO	0.93	0.94	0.93	0.93	0.93
PRIVOPs					
PRIVOP/CMD	27.713	27.636	28.804	28.826	28.764
DIAG/CMD	14.801	15.111	21.496	21.217	21.276
DIAG 08/CMD	0.729	0.741	0.724	0.730	0.738
DIAG 10/CMD	5.088	5.305	0.012	0.012	0.012
DIAG 58/CMD	1.238	1.239	1.242	1.242	1.240
DIAG 98/CMD	0.320	0.314	0.312	0.312	0.313
DIAG A4/CMD	2.606	2.608	2.748	2.726	2.745
DIAG A8/CMD	1.635	1.689	1.654	1.684	1.659
DIAG 214/CMD	na	na	11.240	11.547	11.601
SIE/CMD	78.850	78.910	72.313	72.357	72.545
SIE INTCPT/CMD	55.195	54.448	51.343	50.650	50.056
FREE TOTL/CMD	139.504	133.540	102.444	96.476	96.726
VTAM Machines					
WKSET (V)	1585	1529	1435	1402	1446
TOT CPU/CMD (V)	4.4564	4.4557	3.9662	3.9843	3.9795
CP CPU/CMD (V)	2.5686	2.5555	1.9831	1.9921	1.9865
VIRT CPU/CMD (V)	1.8878	1.9002	1.9831	1.9921	1.9930
DIAG 98/CMD (V)	0.320	0.319	0.313	0.316	0.316
SFS Servers					
WKSET (V)	1115	1101	966	997	935
TOT CPU/CMD (V)	4.2722	4.3404	3.7639	3.7171	3.6943
CP CPU/CMD (V)	2.3680	2.4271	1.8268	1.7960	1.7865
VIRT CPU/CMD (V)	1.9043	1.9133	1.9371	1.9211	1.9078
FP REQ/CMD(Q)	1.334	1.337	1.340	1.338	1.339
IO/CMD (Q)	2.003	2.002	1.974	1.982	1.954
IO TIME/CMD (Q)	0.042	0.041	0.037	0.036	0.036
SFS TIME/CMD (Q)	0.131	0.127	0.087	0.088	0.087
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM					

Table 20. Software Mode Comparisons

9221-170 / Minidisk

The measurements in this section examine the performance effects, on VM/ESA 1.1, when changing the user virtual machine mode from 370 mode to XA mode.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9221-170
- STORAGE:
   - RSTOR:   48M
   - XSTOR:   16M (all reserved for MDC)

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  H3AP01     3380
              H3SRV      3380
              H3RES      3380
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
		NUMBER OF PACKS				
3380-A	3 - 3380-03	2	2	2	4	0
3380-D	1 - 3380-03	1	0	1	1	0

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
                  CONTROLLER  NUMBER  LINESPEED
                  3088-02       1       4.5M
```

3) SOFTWARE CONFIGURATION

```
- DRIVER:               TPNS
- THINK TIME DISTR:     BACTRIAN
- CMS BLOCKSIZE:        4K
- USER VM SIZE:         2M
- USER CMS MODE:        370,XA
- USER RELSHARE:        100

- SERVER MACHINES:
                                    VM SIZE/
  SERVER MACHINE   TYPE           CMS MODE   RELSHARE OTHER OPTIONS

  VTAM             VTAM/VSCS      64M/XA     10000    QUICKDSP ON
```

4) MEASUREMENT DISCUSSION

Comparing the 370 mode users to the XA mode users, internal throughput decreased by 2.8% for this environment. This was due to emulation per command (EMUL/CMD(H)) increased by 5.8%. The increase was larger than the increase experienced on the 9021-720 since the unique instructions needed to implement XA mode do not perform as well on the 9221-170 (see section “Software Mode Comparisons” on page 119 for more details). External response time (AVG LAST(T)) increased by 14.6%.

USER MODE RELEASE RUN ID	370 ESA 1.1 H14R0287	XA ESA 1.1 H14R0289
Environment		
REAL STORAGE	**48M**	**48M**
EXP. STORAGE	**16M**	**16M**
USERS	**280**	**280**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
Response Time		
TRIV INT	0.172	0.171
NONTRIV INT	0.810	0.947
TOT INT	0.624	0.714
TOT INT ADJ	0.529	0.607
AVG FIRST (T)	0.270	0.280
AVG LAST (T)	0.550	0.630
Throughput		
AVG THINK (T)	28.39	28.31
ETR	8.23	8.26
ETR (T)	9.71	9.71
ETR RATIO	0.848	0.851
ITR (H)	11.90	11.57
ITR	10.10	9.84
EMUL ITR	15.99	15.20
ITRR (H)	1.000	0.972
ITRR	1.000	0.973
Proc. Usage		
PBT/CMD (H)	84.060	86.445
PBT/CMD	84.486	86.495
CP/CMD (H)	36.380	36.005
CP/CMD	30.910	30.891
EMUL/CMD (H)	47.680	50.440
EMUL/CMD	53.577	55.604
Processor Util.		
TOTAL (H)	81.59	83.95
TOTAL	82.00	84.00
UTIL/PROC (H)	81.59	83.95
UTIL/PROC	82.00	84.00
TVR(H)	1.76	1.71
TVR	1.58	1.56
Storage		
WKSET (V)	77	75
PGBLPGS	9520	9516
PGBLPGS/USER	34.0	34.0
FREEPGS	804	802
FREE UTIL	0.88	0.88
SHRPGS	900	907
Paging		
READS/SEC	62	62
WRITES/SEC	49	50
PAGE/CMD	11.437	11.533
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.461	5.663

USER MODE RELEASE RUN ID	370 ESA 1.1 H14R0287	XA ESA 1.1 H14R0289
Environment		
REAL STORAGE	**48M**	**48M**
EXP. STORAGE	**16M**	**16M**
USERS	**280**	**280**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
I/O		
VIO RATE	82	83
VIO/CMD	8.449	8.547
MDC READS	53	53
MDC WRITES	26	26
MDC MODS	21	21
MDC HIT RATIO	0.91	0.91
PRIVOPs		
PRIVOP/CMD	14.426	14.615
DIAG/CMD	26.002	25.934
DIAG 08/CMD	0.618	0.618
DIAG 10/CMD	0.000	0.000
DIAG 58/CMD	1.236	1.133
DIAG 98/CMD	2.473	2.368
DIAG A4/CMD	3.812	3.913
DIAG A8/CMD	2.061	2.059
DIAG 214/CMD	11.334	11.842
SIE/CMD	56.977	56.428
SIE INTCPT/CMD	42.163	41.192
FREE TOTL/CMD	97.571	91.438
VTAM Machines		
WKSET (V)	211	299
TOT CPU/CMD (V)	18.3362	18.1798
CP CPU/CMD (V)	8.1785	8.1155
VIRT CPU/CMD (V)	10.1577	10.0643
DIAG 98/CMD (V)	2.519	2.405
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 21. 9221-170 / Minidisk 370 Mode Users Versus XA Mode Users

9221-170 / 35% SFS

The measurements in this section examine the performance effects, on VM/ESA 1.1, when changing the user virtual machine modes from 370 mode to XC mode.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9221-170
- STORAGE:
   - RSTOR:   48M
   - XSTOR:   16M  (all reserved for MDC)

- DASD:
              PACK NAME   TYPE
   - SYSTEM:  H3AP01      3380
              H3SRV       3380
              H3RES       3380
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	3 - 3380-03	2	2	2	4	5
3380-D	1 - 3380-03	1	0	1	1	1

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINESPEED
3088-02	1	4.5M

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR:  BACTRIAN
- CMS BLOCKSIZE:     4K
- USER VM SIZE:      2M
- USER CMS MODE:     370,XC
- USER RELSHARE:     100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON
RWSERVE1 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
CRRRECOV	CRR	32M/XC	100	NONE

4) MEASUREMENT DISCUSSION

Comparing the 370 mode users to the XC mode users, internal throughput decreased by 2.1% for this environment. This was due to emulation per command (EMUL/CMD(H)) increased by 5.6%. This was consistent with the results on the 9021-720 (see section "Software Mode Comparisons" on page 119). External response time (AVG LAST(T)) increased by 2.9%.

USER MODE RELEASE RUN ID	370 ESA 1.1 H14F0241	XC ESA 1.1 H14F0242
Environment		
REAL STORAGE	**48M**	**48M**
EXP. STORAGE	**16M**	**16M**
USERS	**240**	**240**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
Response Time		
TRIV INT	0.184	0.176
NONTRIV INT	1.034	1.037
TOT INT	0.789	0.782
TOT INT ADJ	0.668	0.674
AVG FIRST (T)	0.300	0.310
AVG LAST (T)	0.690	0.710
Throughput		
AVG THINK (T)	28.05	27.90
ETR	7.11	7.27
ETR (T)	8.40	8.43
ETR RATIO	0.847	0.862
ITR (H)	10.15	9.94
ITR	8.59	8.57
EMUL ITR	14.07	13.70
ITRR (H)	1.000	0.979
ITRR	1.000	0.998
Proc. Usage		
PBT/CMD (H)	98.483	100.610
PBT/CMD	98.853	100.778
CP/CMD (H)	44.457	43.587
CP/CMD	38.112	37.940
EMUL/CMD (H)	54.025	57.023
EMUL/CMD	60.741	62.838
Processor Util.		
TOTAL (H)	82.69	84.86
TOTAL	83.00	85.00
UTIL/PROC (H)	82.69	84.86
UTIL/PROC	83.00	85.00
TVR(H)	1.82	1.76
TVR	1.63	1.60
Storage		
WKSET (V)	82	81
PGBLPGS	9623	9638
PGBLPGS/USER	40.1	40.2
FREEPGS	730	732
FREE UTIL	0.87	0.87
SHRPGS	1080	1091
Paging		
READS/SEC	58	57
WRITES/SEC	47	46
PAGE/CMD	12.505	12.212
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.717	5.572

USER MODE RELEASE RUN ID	370 ESA 1.1 H14F0241	XC ESA 1.1 H14F0242
Environment		
REAL STORAGE	**48M**	**48M**
EXP. STORAGE	**16M**	**16M**
USERS	**240**	**240**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
I/O		
VIO RATE	56	57
VIO/CMD	6.670	6.758
MDC READS	47	47
MDC WRITES	16	16
MDC MODS	11	11
MDC HIT RATIO	0.89	0.89
PRIVOPs		
PRIVOP/CMD	22.498	22.734
DIAG/CMD	23.772	23.060
DIAG 08/CMD	0.715	0.711
DIAG 10/CMD	0.000	0.000
DIAG 58/CMD	1.191	1.186
DIAG 98/CMD	2.620	2.371
DIAG A4/CMD	2.501	2.490
DIAG A8/CMD	1.667	1.778
DIAG 214/CMD	10.838	10.908
SIE/CMD	69.554	67.936
SIE INTCPT/CMD	53.557	51.632
FREE TOTL/CMD	111.239	104.809
VTAM Machines		
WKSET (V)	194	190
TOT CPU/CMD (V)	18.9415	18.2532
CP CPU/CMD (V)	8.5119	8.2056
VIRT CPU/CMD (V)	10.4296	10.0477
DIAG 98/CMD (V)	2.710	2.480
SFS Servers		
WKSET (V)	401	397
TOT CPU/CMD (V)	12.7847	12.9280
CP CPU/CMD (V)	6.8970	6.9999
VIRT CPU/CMD (V)	5.8877	5.9281
FP REQ/CMD(Q)	1.340	1.339
IO/CMD (Q)	2.105	2.144
IO TIME/CMD (Q)	0.036	0.037
SFS TIME/CMD (Q)	0.051	0.050

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

Table 22. 9221-170 35% SFS Comparing 370 Mode Users to XC Mode Users

OfficeVision Migration from VM/XA 2.1

9021-720

The following 9021-720 runs are provided to show the effects of VM/ESA 1.0 and VM/ESA 1.1 on an OfficeVision environment.

1) WORKLOAD: IOB V2.1

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-720
- STORAGE:
   - RSTOR:   512 M
   - XSTOR:   2048 M

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              PSPT01     3380-D
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-D	20 - 3880-3	20	16	12	40	0
3380-D	1 - 3880-G23	0	0	0	0	4
3380-A	3 - 3880-G23	0	0	0	0	12
3380-K	2 - 3990-2	0	0	0	0	16

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	36	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: IOB
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:    100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VM/XA & VTAM 3.2:				
VTAMA	VTAM/VSCS	16M/370	10000	QUICKDSP ON
VTAMB	VTAM/VSCS	16M/370	10000	QUICKDSP ON
VTAMC	VTAM/VSCS	16M/370	10000	QUICKDSP ON
VSCS2	VSCS	8M/370	10000	QUICKDSP ON
VSCS4	VSCS	8M/370	10000	QUICKDSP ON
VSCS6	VSCS	8M/370	10000	QUICKDSP ON
PRODBM	OV/VM	16M/XA	10000	QUICKDSP ON
PROCAL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMAIL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMBX00 - 50	OV/VM	16M/XA	10000	QUICKDSP ON; IBCENTRL=Y
VM/ESA & VTAM 3.3:				
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
PRODBM	OV/VM	16M/XA	10000	QUICKDSP ON
PROCAL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMAIL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMBX00 - 50	OV/VM	16M/XA	10000	QUICKDSP ON; IBCENTRL=Y

4) MEASUREMENT DISCUSSION

This section documents the migration data collected for an OV/VM environment. The base starting point was 6000 users running on VM/XA 2.1. With the hardware configuration available, adequate performance was achieved at about 87% CPU utilization and an external response time (AVG LAST (T)) of 0.98 seconds. A measurement was made increasing the users to 6200. This resulted in a very large increase in external response time (109%) and a reduction in the internal throughput rate (ITR (H)), indicating that the system had become over loaded.

Using the 6000 user VM/XA 2.1 measurement as the base, VM was upgraded to VM/ESA 1.0. With this environment a positive effect was observed on external response time, improving by 0.170 seconds (17%) with a slight decrease in the internal throughput rate. Again, a measurement was made increasing the number of users to 6200. This time the external response time only increased to 1.05 seconds (30%) and the internal throughput rate remained about the same, indicating that VM/ESA 1.0 could support this increased user load.

Using the 6200 user VM/ESA 1.0 measurement as the base, VM was upgraded to VM/ESA 1.1. Aan additional positive effect was observed on external response time, improving by 0.126 seconds (12%) with an ITR (H) improvement of about 4.1%.

In summary, both VM/ESA 1.0 and VM/ESA 1.1 have improved the performance of this OfficeVision environment.

RELEASE RUN ID	XA SP 2.1 Y6$V6002	XA SP 2.1 Y6$V6201	ESA 1.0 Y63V6001	ESA 1.0 Y63V6203	ESA 1.1 Y64V620F
Environment					
REAL STORAGE	**512M**	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**	**2048M**
USERS	**5998**	**6200**	**6000**	**6200**	**6201**
VTAMs	**3**	**3**	**1**	**1**	**1**
VSCSs	**3**	**3**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**	**6**
Response Time					
TRIV INT	0.046	0.068	0.037	0.037	0.036
NONTRIV INT	0.404	0.521	0.381	0.407	0.403
TOT INT	0.314	0.375	0.284	0.292	0.303
TOT INT ADJ	0.357	0.481	0.356	0.378	0.375
AVG FIRST (T)	0.743	1.640	0.603	0.803	0.627
AVG LAST (T)	0.980	2.050	0.810	1.053	0.927
Throughput					
AVG THINK (T)	41.97	42.02	42.58	42.62	42.50
ETR	127.04	145.92	140.21	148.93	142.20
ETR (T)	111.66	113.86	111.70	115.11	115.02
ETR RATIO	1.138	1.282	1.255	1.294	1.236
ITR (H)	128.63	123.73	125.81	125.78	130.96
ITR	25.25	27.41	26.37	27.54	27.05
EMUL ITR	47.74	52.98	51.34	53.40	50.20
ITRR (H)	1.000	0.962	0.978	0.978	1.018
ITRR	1.000	1.086	1.044	1.091	1.071
Proc. Usage					
PBT/CMD (H)	46.646	48.493	47.690	47.701	45.817
PBT/CMD	46.661	48.394	47.716	47.084	45.731
CP/CMD (H)	23.787	25.309	24.279	24.286	22.595
CP/CMD	21.942	23.363	23.187	22.847	21.040
EMUL/CMD (H)	22.854	23.180	23.405	23.410	23.217
EMUL/CMD	24.718	25.031	24.530	24.237	24.691
Processor Util.					
TOTAL (H)	520.84	552.13	532.70	549.10	526.99
TOTAL	521.00	551.00	533.00	542.00	526.00
UTIL/PROC (H)	86.81	92.02	88.78	91.52	87.83
UTIL/PROC	86.83	91.83	88.83	90.33	87.67
TVR(H)	2.04	2.09	2.04	2.04	1.97
TVR	1.89	1.93	1.95	1.94	1.85
Storage					
WKSET (V)	49	47	55	55	53
PGBLPGS	90929	89983	109K	109K	108K
PGBLPGS/USER	15.2	14.5	18.6	18.0	17.8
FREEPGS	12142	12264	11886	12315	13129
FREE UTIL	0.97	0.98	0.97	0.96	0.97
SHRPGS	1161	1190	1369	1350	1349
Paging					
READS/SEC	0	0	53	56	34
WRITES/SEC	24	30	66	82	77
PAGE/CMD	0.215	0.263	1.065	1.199	0.965
XSTOR IN/SEC	1824	2125	1442	1538	1640
XSTOR OUT/SEC	1981	2334	1638	1746	1796
XSTOR/CMD	34.077	39.163	27.573	28.528	29.873
FAST CLR/CMD	29.671	29.888	24.315	23.968	14.980

RELEASE RUN ID	XA SP 2.1 Y6$V6002	XA SP 2.1 Y6$V6201	ESA 1.0 Y63V6001	ESA 1.0 Y63V6203	ESA 1.1 Y64V620F
Environment					
REAL STORAGE	**512M**	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**	**2048M**
USERS	**5998**	**6200**	**6000**	**6200**	**6201**
VTAMs	**3**	**3**	**1**	**1**	**1**
VSCSs	**3**	**3**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**	**6**
I/O					
VIO RATE	2337	2427	2288	2316	2498
VIO/CMD	20.930	21.316	20.483	20.119	21.718
MDC READS	1027	1058	1141	1155	1369
MDC WRITES	931	968	978	991	979
MDC MODS	828	859	837	847	833
MDC HIT RATIO	0.90	0.90	0.89	0.88	0.90
PRIVOPs					
PRIVOP/CMD	20.057	19.605	23.070	21.660	22.684
DIAG/CMD	57.629	58.677	63.878	62.669	85.547
DIAG 08/CMD	8.732	9.082	9.516	9.226	9.372
DIAG 10/CMD	13.765	13.851	14.745	14.516	0.000
DIAG 58/CMD	1.791	1.792	1.862	1.842	1.861
DIAG 98/CMD	0.394	0.369	0.412	0.408	0.409
DIAG A4/CMD	11.311	11.497	10.268	10.086	11.624
DIAG A8/CMD	5.992	6.095	6.338	6.237	6.260
DIAG 214/CMD	na	na	na	na	35.185
SIE/CMD	146.734	152.894	155.843	151.225	133.542
SIE INTCPT/CMD	89.508	90.207	101.298	96.784	88.138
FREE TOTL/CMD	174.247	179.875	210.847	204.599	222.570
VTAM Machines					
WKSET (V)	429	486	655	1823	1869
TOT CPU/CMD (V)	5.3551	5.3122	5.8015	5.5985	4.9288
CP CPU/CMD (V)	3.0299	3.0370	3.2327	3.0299	2.3847
VIRT CPU/CMD (V)	2.3252	2.2752	2.5688	2.5686	2.5440
DIAG 98/CMD (V)	0.409	0.392	0.420	0.418	0.414
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM					

Table 23. The Effects of Migrating an OfficeVision Environment to VM/ESA 1.1

MVS Guest Migration from VM/ESA 1.0

3090-600J (1 CPU)

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

1) WORKLOAD: CB84

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  3090-600J   2 CPUs configured; 1 dedicated
- STORAGE:
   - RSTOR:   256M
   - XSTOR:   512M

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380
              DRV308     3380
              USERPK     3380
              PPLOAD     3380
              PAGEPK     3380
              PAGEF1     3380
              PROGPK     3380
              POOLS1     3380
              POOLS2-6   3380
              POOLS7     3380
              POOLS8     3380
              POOLS9     3380
              CB8413     3380
              CB84LB     3380
              CBLOAD     3380
              CKPTDS     3380
              STGF31-37  3380
              STG637     3380
              IDAVOL     3380
              PPL664     3380
              PPL665     3380


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

- TAPE:       MONITOR 3480
```

3) SOFTWARE CONFIGURATION

```
- MVS VERSION:      3.1.0e
- GUEST VM SIZE:    256M
- V=R SIZE:         212M
```

4) MEASUREMENT DISCUSSION

The following table contains the measurement data for the CB84 runs that were made with native MVS, MVS as a guest of VM/ESA 1.0, and MVS as a guest of an early version of VM/ESA 1.1. (We have no reason to believe later versions of VM/ESA 1.1 would behave differently).

V=R guest operating system performance on VM/ESA 1.1 was predicted to be equivalent to VM/ESA 1.0 because no guest performance improvements for V=R guests were made to VM/ESA 1.1. The results of these measurements, as detailed in the following table, show that was the case.

GUEST TYPE VM RELEASE RUN ID	**native** **n/a** **CBNTV310**	**V=R** **VM/ESA 1.0** **CBES10VR**	**V=R** **VM/ESA 1.1** **CBES11VR**
Configuration			
Num. Processors	1	1 dedicated	1 dedicated
Real Storage	256M	256M	256M
Expanded Storage	512M	512M	512M
Throughput			
Int Throughput (ITR)	0.478	0.454	0.455
ITR % of Native	100	95.0	95.2
Ext Throughput (ETR)	0.467	0.444	0.443
ETR % of Native	100	95.1	94.9
Processor Data			
Elapsed Seconds	835.841	878.253	880.463
Processor Busy %	97.59	97.79	97.42
Processor Seconds	815.697	858.844	857.747
Batch Data			
Num. of Initiators	24	23	23
Num. of Batch Jobs	390	390	390
Chan Path/DASD Data			
No. of DASD Paths	32	32	32
Avg. Ch Path Busy	8.86	9.07	8.84
High Ch Path Busy	24.86	25.32	24.39
I/O Interrupt Rate	575.70	547.50	546.20
I/O Interrupts/Tran	1233.83	1232.93	1233.10
Paging			
Total: In+Out	0.00	0.00	0.00
NSW/NVIO: Total	0.00	0.00	0.00
VIO: Total	0.00	0.00	0.00
Main Storage Data			
Storage Size	256M	128M	128M
Ext. Storage Size	512M	256M	256M
Total Frames			
SQA - Avg	4452	841	843
LPA - Avg	697	702	701
CSA - Avg	278	277	277
LSQA - Avg	1089	1068	1065
Priv Area - Avg	7282	7230	6993
Unused - Avg	49610	21738	21977
Total - Avg	65086	32766	32766
Fixed Frames			
SQA - Avg	4419	813	815
LPA - Avg	32	32	33
CSA - Avg	0	0	0
LSQA - Avg	962	946	942
Priv Area - Avg	945	958	817
Below 16M - Avg	287	299	261
Nucleus	1675	907	907
Tot Fixed - Avg	8038	3658	3508
Exp. Storage Frames			
SQA - Avg	0	0	0
LPA - Avg	0	0	0
CSA - Avg	0	0	0
LSQA - Avg	5	6	6
Priv Area - Avg	498	455	474
Unused - Avg	130519	65029	65084
Total - Avg	131072	65536	65536

Table 24. CB84 Measurement Data: VM/ESA 1.0 And VM/ESA 1.1 MVS 3.1.0e Guests

VSE Guest Migration from VM/ESA 1.0

VSE configuration

The configuration for VSE guest measurements was as follows:

- Partition priority was BG=FB=FA=F9=F8=F7=F6=F5=F4,F2,F3,F1. The important point is that F4 through FB were of equal priority.
- The job classes for F4 through FB were configured to mix and balance the workload across the partitions.

The following MAP command output shows the VSE system storage allocation:

```
SPACE AREA   PRTY   V-SIZE GETVIS V-ADDR    R-SIZE R-ADDR NAME
  S   SUP             844K             0      780K      0 $$A$SUPX
  S   SVA            1636K  1360K  D3000    52276K
  S   UNUSED          256K        3C0000
  1   BG V     4     1500K   804K 400000      144K 440000 NO NAME
  1   F1 V     1      768K   256K 640000      256K 400000 POWSTART
  1   F4       4      500K   524K 740000      144K 4F2000 NO NAME
  2   FB       4      500K   524K 840000      144K 5EE000 NO NAME
  2   UNUSED         6912K        940000
  2   F2 V     3     5900K  5364K 400000      144K 464000 CICSICCF
  2   UNUSED         1024K        F00000
  3   F3 V     2     3300K  4828K 400000      424K 488000 VTAMSTRT
  3   UNUSED         4160K        BF0000
  4   F5 V     4      500K   524K 400000      144K 516000 NO NAME
  4   F6 V     4      500K   524K 500000      144K 53A000 NO NAME
  4   UNUSED        10240K        600000
  5   F7 V     4      500K   524K 400000      144K 55E000 NO NAME
  5   F8 V     4      500K   524K 500000      144K 582000 NO NAME
  5   UNUSED        10240K        600000
  6   F9 V     4      500K   524K 400000      144K 5A6000 NO NAME
  6   FA V     4      500K   524K 500000      144K 5CA000 NO NAME
  6   UNUSED        10240K        600000
      AVAIL         47168K                  11408K
      TOTAL         81920K                  16384K
```

Although VTAM and CICS/ICCF appear in this particular MAP output, they were always shut down before a measurement.

The storage configuration for the MODE=VMESA (single address space) runs was the same for the F4 through FB partition allocation sizes. Except F1, the other partition sizes were reduced but were not used during the measurements.

The LST queue was deleted before each run; if the accounting file was large, it was deleted as well.

Overview

All data for VSE guest performance shown here represent the PACEX8 VSE batch workload; a CICS interactive workload would show different results. An above-average I/O rate (compared to a *typical* commercial batch environment) is one characteristic of the PACEX8 workload; keep this in mind when examining the data.

Brief observations precede each of the following tables. The first three tables show how the various VSE modes compare across VM/XA 2.1, VM/ESA 1.0 ESA Feature, and VM/ESA 1.1. As expected, these VM releases performed similarly; no changes in these releases have had any significant effect on VSE guest performance.

Two tables follow that show performance of VSE guests of VM/ESA 1.1 running in BASIC mode and then, with the same configurations, in a dedicated LPAR.

The final VSE guest table demonstrates the dramatic effect that CCW fast path can have. This enhancement is available as APAR VM51012 for VM/ESA 1.1 systems. This capability was not present in any of the systems used for measurements summarized in the previous tables.

Some values shown in the tables require some explanation:

ITR is calculated as the number of batch jobs (56) divided by the number of CPU busy minutes. The CPU busy minutes is calculated as the elapsed time for completion of the batch jobs multiplied by the CPU busy percent (from the VMPRF report) divided by 100.

ETR is calculated as the number of batch jobs (56) divided by the number of minutes for completion.

DASD I/O per second is calculated as the number of DASD reads/writes (from the VMPRF report) divided by the number of elapsed seconds for batch job completion.

elapsed seconds is the total number of seconds required for completion of all 56 batch jobs comprising the PACEX8 workload.

processor seconds is elapsed seconds multiplied by processor busy % divided by 100--this is also known as CPU busy time.

Other values in the tables were taken directly from the VMPRF reports.

VM Release Comparisons / VSE V=V MODE=ESA Guest

Unless otherwise noted, all of the VSE measurements discussed in this report used the following configuration.

1) WORKLOAD: PACEX8

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9121-320
- STORAGE:
   - RSTOR:   128M
   - XSTOR:   0M

- DASD:
              PACK NAME  TYPE  CHANNEL PATHS
   - SYSTEM:  PRFRES     3380        2
              PRF01      3380        2
              DOSRES     3380        2
              SYSWK1     3380        2

      TYPE OF    NUMBER/TYPE OF
       DASD      CONTROL UNIT     NUMBER OF PACKS
      3380-D       2 - 3880-05     2 VM system
      3380-D       2 - 3880-03     2 VSE system
      3380-A       4 - 3880-03    10 VSAM data

- TAPE:       MONITOR  3480
```

3) SOFTWARE CONFIGURATION

```
- GUEST VM SIZE:      16M
- GUEST MACHINE MODE: ESA
```

4) MEASUREMENT DISCUSSION

Table 25 on page 138 is the first of three tables showing comparisons of the various modes of VSE guests of VM. This case shows MODE=ESA (multi-address space) VSE running as a V=V guest on each of three VM releases: VM/XA 2.1, VM/ESA 1.0 ESA Feature, and VM/ESA 1.1.

The ITRs for the workload on the three releases were similar; differences, considered within measurement tolerances, were not significant. Other indicators, including ETR and T/V ratios, were all similar, confirming that, from a VSE guest perspective, there is little difference across these three VM releases.

RELEASE RUN ID	**VM/XA 2.1** **PD10608C**	**VM 1.0 ESA** **PD10606A**	**VM/ESA 1.1** **PD10604A**
Environment			
processor model	9121-320	9121-320	9121-320
real storage	128	128	128
expanded storage	0	0	0
virtual machine type	V=V	V=V	V=V
SET MACHINE	ESA	ESA	ESA
VSE release	VSE/ESA 1.1.0	VSE/ESA 1.1.0	VSE/ESA 1.1.0
VSE MODE	MODE=ESA	MODE=ESA	MODE=ESA
Throughput			
ITR	14.930	15.009	14.680
ITRR	1.00	1.01	0.98
ETR	6.211	6.154	5.989
ETRR	1.00	0.99	0.96
Processor Data			
elapsed seconds	541	546	561
processor busy %	41.6	41.0	40.8
processor seconds	225.1	223.9	228.9
T/V ratio	1.72	1.70	1.71
Paging			
reads/second	0	0	0
writes/second	0	0	0
I/O			
DASD reads & writes	159121	169847	169847
DASD I/O per second	294.1	311.1	302.8

Table 25. VSE V=V MODE=ESA Guest

VM Release Comparisons / VSE V=V MODE=VMESA Guest

The observations made previously about the MODE=ESA case apply equally well in the MODE=VMESA case as shown in Table 26. Again, the ITRs were similar except for the VM/ESA 1.0 ITR; its ITR appeared lower, even considering the measurement tolerance, but not by much.

The other major indicators confirmed the similarity between the performance of the VM releases for this case.

RELEASE RUN ID	**VM/XA 2.1** **PD10608D**	**VM 1.0 ESA** **PD10606E**	**VM/ESA 1.1** **PD10606F**
Environment			
processor model	9121-320	9121-320	9121-320
real storage	128	128	128
expanded storage	0	0	0
virtual machine type	V=V	V=V	V=V
SET MACHINE	ESA	ESA	ESA
VSE release	VSE/ESA 1.1.0	VSE/ESA 1.1.0	VSE/ESA 1.1.0
VSE MODE	MODE=VMESA	MODE=VMESA	MODE=VMESA
Throughput			
ITR	18.174	17.204	18.107
ITRR	1.00	0.95	1.00
ETR	6.234	6.022	6.211
ETRR	1.00	0.97	1.00
Processor Data			
elapsed seconds	539	558	541
processor busy %	34.3	35.0	34.3
processor seconds	184.9	195.3	185.6
T/V ratio	1.82	1.84	1.80
Paging			
reads/second	0	0	0
writes/second	0	0	0
I/O			
DASD reads & writes	158793	169355	169355
DASD I/O per second	294.6	303.5	313.0

Table 26. VSE V=V MODE=VMESA Guest

VM Release Comparisons / VSE V=R MODE=ESA Guest

The V=R MODE=ESA comparison shown in Table 27 shows nearly identical performance across the three VM releases.

RELEASE RUN ID	**VM/XA 2.1** **PD10608E**	**VM 1.0 ESA** **PD10606B**	**VM/ESA 1.1** **PD10605A**
Environment			
processor model	9121-320	9121-320	9121-320
real storage	128	128	128
expanded storage	0	0	0
virtual machine type	V=R	V=R	V=R
SET MACHINE	ESA	ESA	ESA
VSE release	VSE/ESA 1.1.0	VSE/ESA 1.1.0	VSE/ESA 1.1.0
VSE MODE	MODE=ESA	MODE=ESA	MODE=ESA
Throughput			
ITR	26.187	25.809	26.332
ITRR	1.00	0.99	1.01
ETR	6.154	6.143	6.109
ETRR	1.00	1.00	0.99
Processor Data			
elapsed seconds	546	547	550
processor busy %	23.5	23.8	23.2
processor seconds	128.3	130.2	127.6
T/V ratio	1.10	1.11	1.09
Paging			
reads/second	0	0	0
writes/second	0	0	0
I/O			
DASD reads & writes	159021	169770	169770
DASD I/O per second	291.3	310.4	308.7

Table 27. VSE V=R MODE=ESA Guest

VM/ESA 1.1 / VSE Guest MODE Comparison

VM/ESA 1.1 behaved as expected as a host for a batch VSE system on a 9121 processor. The MODE=VMESA V=V VSE attained higher throughput than the MODE=ESA V=V case due to the full VM handshaking used by the MODE=VMESA guest. The V=R case performed dramatically better than either of the two previous configurations. This was due, for the most part, to the I/O benefits provided in the V=R case: IOASSIST was ON and CCWTRAN was OFF (i.e., no CCW translation is performed). "VM/ESA 1.1 in LPAR / VSE Guest MODE Comparison" on page 142 shows that these were the major contributing benefits for V=R guests.

RELEASE RUN ID	**VM/ESA 1.1** **PD10604A**	**VM/ESA 1.1** **PD10606F**	**VM/ESA 1.1** **PD10605A**
Environment			
processor model	9121-320	9121-320	9121-320
real storage	128	128	128
expanded storage	0	0	0
virtual machine type	V=V	V=V	V=R
SET MACHINE	ESA	ESA	ESA
VSE release	VSE/ESA 1.1.0	VSE/ESA 1.1.0	VSE/ESA 1.1.0
VSE MODE	MODE=ESA	MODE=VMESA	MODE=ESA
Throughput			
ITR	14.680	18.107	26.332
ITRR	1.00	1.23	1.79
ETR	5.989	6.211	6.109
ETRR	1.00	1.04	1.02
Processor Data			
elapsed seconds	561	541	550
processor busy %	40.8	34.3	23.2
processor seconds	228.9	185.6	127.6
T/V ratio	1.71	1.80	1.09
Paging			
reads/second	0	0	0
writes/second	0	0	0
I/O			
DASD reads & writes	169644	170506	169698
DASD I/O per second	302.4	315.2	308.5

Table 28. VSE as Guest of VM/ESA 1.1

VM/ESA 1.1 in LPAR / VSE Guest MODE Comparison

Previous tables have shown the similarities in the performance of VSE guests across the three most recent VM releases. Although not shown here, measurements confirmed that the same pattern applies to the same configurations placed in an LPAR. That is, it can be shown that V=V MODE=ESA VSE guests of the three VM releases discussed here will behave in a similar manner in an LPAR, as will the other VSE configurations.

To point out the major differences that occur in an LPAR, Table 29 on page 143 shows the performance characteristics of the VSE guest configurations running in a dedicated LPAR. This is not a recommended configuration for a single-processor environment, but it serves to illustrate some key points. Ordinarily, shared LPARs would be used in a uniprocessor environment. One or more dedicated LPARs may make sense on 9121-480 and larger processors (with 2 or more CPUs).

The MODE=VMESA ITR improved 20 percent over the MODE=ESA V=V ITR (similar to the ratio between the same two configurations in the non-LPAR case shown in Table 28 on page 141). The V=R configuration, however, did not show the significantly improved ITR seen in the non-LPAR case. The ITR for the V=R configuration fell in a neighborhood between the two V=V configurations. This is primarily because IOASSIST was OFF and CCWTRAN was ON in the V=R LPAR environment. Thus we see the magnitude of the benefit of these features when we can no longer use them because of the LPAR environment.

Although this table shows the LPAR characteristics for VM/ESA 1.1, other measurements have shown that similar patterns hold for the other two VM releases as well.

Avoid the temptation to directly compare this table with Table 28 on page 141 in an attempt to do a head-to-head BASIC versus LPAR mode comparison. VMPRF data used for the tables does not include LPAR overhead. It is known to be around three to five percent of the CPU resource for a dedicated LPAR.

RELEASE RUN ID	**VM/ESA 1.1** **PD10605C**	**VM/ESA 1.1** **PD10606G**	**VM/ESA 1.1** **PD10605B**
Environment			
processor model	9121-320	9121-320	9121-320
real storage	122	122	122
expanded storage	0	0	0
virtual machine type	V=V	V=V	V=R
SET MACHINE	ESA	ESA	ESA
VSE release	VSE/ESA 1.1.0	VSE/ESA 1.1.0	VSE/ESA 1.1.0
VSE MODE	MODE=ESA	MODE=VMESA	MODE=ESA
Throughput			
ITR	14.278	17.142	15.094
ITRR	1.00	1.20	1.06
ETR	6.054	6.154	6.098
ETRR	1.00	1.02	1.01
Processor Data			
elapsed seconds	555	546	551
processor busy %	42.4	35.9	40.4
processor seconds	235.3	196.0	222.6
T/V ratio	1.72	1.83	1.68
Paging			
reads/second	0	0	0
writes/second	0	0	0
I/O			
DASD reads & writes	170216	170164	170590
DASD I/O per second	306.7	311.7	309.6

Table 29. VSE Guest of VM/ESA 1.1 in LPAR

VM/ESA 1.1 CCW Fast Path Benefit for VSE

The following section involves measurements made on a different configuration than those discussed previously.

1) WORKLOAD: PACEX8

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9121-480 CPU 1 only
- STORAGE:
   - RSTOR:   256M
   - XSTOR:   0M

- DASD:
              PACK NAME  TYPE  CHANNEL PATHS
   - SYSTEM:  PSPT01     3380        2
              PSYS02     3380        2
              DOSRES     3380        2
              SYSWK1     3380        2

      TYPE OF    NUMBER/TYPE OF
       DASD      CONTROL UNIT     NUMBER OF PACKS
      3380-D       2 - 3880-23     4 VM and VSE system packs
      3380-A       2 - 3880-03    10 VSE VSAM data

- TAPE:       MONITOR  3480
```

3) SOFTWARE CONFIGURATION

```
- GUEST VM SIZE:      16M
- GUEST MACHINE MODE: ESA
```

4) MEASUREMENT DISCUSSION

As mentioned earlier, APAR VM51012 for VM/ESA 1.1 introduces an enhanced CCW translation capability that is particularly useful for V=V VSE guests (of any MODE). The CCW fast path enhancement enables VM to do a low-overhead translation for simple DASD channel programs. VM translates more complex or unconventional channel programs in the usual manner, requiring the increased overhead. (See “CCW Fast Path” on page 9 for a more detailed explanation.) The enhanced capability handled approximately 97% of the DASD channel programs used in the VSE batch workload, enabling dramatic throughput improvements.

Table 30 on page 145 represents some measurements to show two different cases: in one case the VSE VSAM data resided on dedicated 3380 volumes and in the other, the VSE VSAM data resided on 3380 minidisks. For each of these two cases, measurements were made with and without the benefit of the CCW fast path code. In all cases, the VSE was running as a V=V MODE=ESA guest of VM/ESA 1.1.

For the dedicated DASD case we saw an improvement of almost 28% in the ITR when CCW fast path was used. For the minidisk environment we saw an improvement of nearly 37% with CCW fast path.

The VSE PACEX8 workload magnified the benefit of the CCW fast path capability since the workload is very I/O intensive. The benefit also depends upon the amount of CPU resource expended for each I/O operation as well as the complexity of the CCW programs. The benefit achieved in a particular instance can be estimated as follows: For the dedicated DASD case, data from the table indicate that a savings of 0.31 millisecond of CPU time per DASD I/O is achieved. For the minidisk case the savings is 0.41 millisecond per DASD I/O. For a selected interval of time, multiply the total number of DASD I/Os by the savings achieved per I/O. Subtract this amount from the CPU seconds used during the interval. This yields an approximation of the CPU seconds used for that interval when CCW fast path is used. From there, a new ITR can be projected.

Using the dedicated DASD data in the table as an example, the initial case shows a total of 162810 DASD I/Os. Multiply this by 0.31 milliseconds per I/O for 50.47 CPU seconds savings. During this interval, project that 225.5 - 50.47 = 175 CPU seconds will be consumed. This is close to the measured value of 176.8. The projected ITR would be the current ITR (14.901) times 1.29 (225.5/175) for a projected ITR of 19.2, once again close to the measured value of 19.006 in this case.

Note that the 0.31 and 0.41 millisecond values for dedicated and minidisk environments apply only to the 9121 processor and for the type of DASD I/O used in the PACEX8 workload. To make a projection for other situations, find the ratio of ITRs for the same workload on the 9121 and the target processor. Apply this ratio to the 0.31 and 0.41 numbers to find the appropriate values for the target processor and then proceed as described above to calculate the CPU time adjustment. Be sure to keep in mind that this is only an estimate of the change.

VSAM DATA SETS	**Dedicated DASD**		**Minidisks**	
CCW FAST PATH RUN ID	**Absent** **PD11030C**	**Present** **PD11104A**	**Absent** **PD11030B**	**Present** **PD11030A**
Environment				
processor model	9121-320	9121-320	9121-320	9121-320
real storage	256	256	256	256
expanded storage	0	0	0	0
virtual machine type	V=V	V=V	V=V	V=V
SET MACHINE	ESA	ESA	ESA	ESA
VSE release	VSE/ESA 1.1.0	VSE/ESA 1.1.0	VSE/ESA 1.1.0	VSE/ESA 1.1.0
VSE MODE	MODE=ESA	MODE=ESA	MODE=ESA	MODE=ESA
Throughput				
ITR	14.901	19.006	13.906	18.991
ITRR	1.00	1.28	1.00	1.37
ETR	7.257	7.336	7.134	7.368
ETRR	1.00	1.01	1.00	1.03
Processor Data				
elapsed seconds	463	458	471	456
processor busy %	48.7	38.6	51.3	38.8
processor seconds	225.5	176.8	241.6	176.9
T/V ratio	1.77	1.38	1.89	1.40
Paging				
reads/second	0	0	0	0
writes/second	0	0	0	0
I/O				
DASD reads & writes	162810	164504	161472	164340
DASD I/O per second	351.6	359.2	342.8	360.4

Table 30. CCW Fast Path Benefit for VSE Guest

8. Hardware Capacity

Processor Capacity

9021 / Minidisk

The processor capacity line measurements were made using a 9021-720 processor. For each measurement the 9021-720 (6-way) was configured for the appropriate storage size and any excess processors were varied offline. The processor utilization of the existing 6-way measurement was chosen as the target for these measurements. The RETAIN XSTOR MDC tuning parameter was set at 32M for the 1-way and 2-way runs and at 64M for the 3-way and 6-way ones. The maximum value for MDC was the amount of available XSTOR.

The 1-way and 2-way measurements had a single VTAM with an internal VSCS handling all user traffic. This traffic was driven by one TPNS machine for the 1-way run and two for the 2-way run. The 3-way and 6-way measurements had a single VTAM with an unused internal VSCS and two external VSCS virtual machines through which all users were connected. Both had three TPNS machines driving the users.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:
   9021-340:       9021-720   CPU 1 only
   9021-500:       9021-720   CPU 0, 1 only
   9021-580:       9021-720   CPU 0, 1, 2 only
   9021-720:       9021-720   CPUs 0-5

- STORAGE:
   9021-340:
     - RSTOR:      128M
     - XSTOR:      256M
   9021-500:
     - RSTOR:      256M
     - XSTOR:      512M
   9021-580:
     - RSTOR:      348M
     - XSTOR:     1024M
   9021-720:
     - RSTOR:      512M
     - XSTOR:     2048M

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              PSPT01     3380-D
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	15 - 3880-3	20	8	12	20	0
3380-D	3 - 3880-3	0	0	0	20	0

```
- TAPE:       MONITOR  3480
```

- COMMUNICATIONS:

	CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
9021-340:	3745-410	1	44	56Kb
9021-500:	3745-410	2	44	56Kb
9021-580:	3745-410	3	44	56Kb
9021-720:	3745-410	3	44	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:             TPNS
- THINK TIME DISTR:   BACTRIAN
- CMS BLOCKSIZE:      4K
- USER VM SIZE:       2M
- USER CMS MODE:      XA
- USER RELSHARE:      100
```

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
9021-340:				
VTAMXAA	VTAM/VSCS	64M/XA	10000	QUICKDSP ON
9021-500:				
VTAMXAA	VTAM/VSCS	64M/XA	10000	QUICKDSP ON
9021-580:				
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
9021-720:				
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON

4) MEASUREMENT DISCUSSION

Changes in Throughput: The main purpose of these measurements was to verify that the performance of VM/ESA Release 1.1 scaled as expected across a selection of the 9021 family. The ITRs ranged from 44.73 for the 1-way to 236.35 for the 6-way with ITRRs of 1, 1.96, 2.76, and 5.28 respectively. These were as expected based on a similar study on the 3090-J family made using VM/ESA Release 1.0 and published in the *VM/ESA Release 1.0 Performance Report*.

The relative drop in ITR/processor as indicated by the ITRRs was caused by

- Normal processing requirements of inter-processor communications
- in the 3-way and 6-way cases, additional overhead generated by using external VSCS virtual machines.

This can be shown by breaking the increase in PBT/CMD into its CP/CMD and EMUL/CMD components. For instance, nearly half of the CP/CMD difference between the 1-way and 3-way measurements can be attributed to the increased cost of the VTAM and VSCS virtual machines, while the rest was caused by normal requirements associated with the increased number of processors. Likewise, 34% of the increase in EMUL/CMD was from additional resource consumption in VTAM and VSCS with the remainder due to inter-processor communications. The CPU usage in the VTAM machines stayed fairly constant between the 1-way and 2-way runs, which both used an internal VSCS, and between the 3-way and 6-way runs, which both had two external VSCS virtual machines.

Other Effects of External VSCS Machines: Likewise, the increase in the total:virtual ratio (TVR (H)) when going from the 1-way and 2-way measurements with an internal VSCS to the 3-way and 6-way measurements was attributable to the use of external VSCS virtual machines.

The 3-way and 6-way measurements showed a large increase in PRIVOP/CMD when compared to the 1-way and 2-way values. About half of this increase was caused by increased IUCV overhead due to the external VSCS virtual machines.

Effects Due to Changes in the Number of Users: When going from the 3-way to the 6-way, external response time (AVE LAST (T)) increased by 0.13 seconds. The 3-way and 6-way runs used identical network configurations, the only difference being that the the 6-way run sent nearly twice as many users through it. Since the internal response time (TOT INT ADJ) changed little between the 3-way and 6-way runs, the increase in external response time can be attributed to an increase in network and VTAM activity in the 6-way run.

In the 6-way run there was a decrease in the amount of storage available per user as compared to the other runs. For the 6-way there were 18.2 pageable pages (PGBLPGS) per user while there were 26.6 for the 3-way measurement. This is because although there was one-third more storage when going from the 3-way to the 6-way, there were nearly twice as many users. The decrease in the available storage per user also influenced the use of XSTOR, causing XSTOR/CMD to jump from 9.686 to 11.414, a 17.8% increase.

Paging was much higher in the 1-way and 2-way runs than it was in the 3-way and 6-way ones. Part of this was due to paging in the shared segments more frequently in the runs with fewer users. As long as the shared segment is active, its storage will not be stolen. However, when there are fewer users, a shared segment will become inactive more often, allowing its frames to be stolen more frequently. In this case, the frames for the 1-way were stolen 3.2 times more frequently than the 6-way. To avoid having key shared segments stolen an installation could keep a user logged on who accesses all of them. There were also fewer pageable pages (PGBLPGS) per user in the 1-way and 2-way runs than the 3-way which also influenced the page rate.

There was also a drop in the number of DIAGNOSE X¢98¢s issued per command in the 6-way run. This reflected an increase in VTAM′s ability to chain I/O buffers for this measurement.

RELEASE RUN ID	ESA 1.1 Y14R1101	ESA 1.1 Y24R2161	ESA 1.1 Y34R3122	ESA 1.1 Y64R5865
Environment				
REAL STORAGE	**128M**	**256M**	**384M**	**512M**
EXP. STORAGE	**256M**	**512M**	**1024M**	**2048M**
USERS	**1100**	**2160**	**3118**	**5860**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**0**	**0**	**2**	**2**
PROCESSORS	**1**	**2**	**3**	**6**
Response Time				
TRIV INT	0.046	0.042	0.045	0.056
NONTRIV INT	0.372	0.306	0.290	0.303
TOT INT	0.252	0.204	0.195	0.193
TOT INT ADJ	0.265	0.217	0.211	0.224
AVG FIRST (T)	0.240	0.235	0.257	0.397
AVG LAST (T)	0.420	0.385	0.397	0.527
Throughput				
AVG THINK (T)	25.66	25.48	25.73	25.19
ETR	41.04	82.21	120.03	238.56
ETR (T)	39.07	77.24	111.09	205.97
ETR RATIO	1.051	1.064	1.080	1.158
ITR (H)	44.73	87.53	123.64	236.35
ITR	46.96	46.56	44.44	45.71
EMUL ITR	67.97	67.53	66.24	69.86
ITRR (H)	1.000	1.957	2.764	5.284
ITRR	1.000	0.991	0.946	0.973
Proc. Usage				
PBT/CMD (H)	22.355	22.850	24.264	25.386
PBT/CMD	22.270	22.916	24.304	25.344
CP/CMD (H)	7.212	7.461	8.395	9.268
CP/CMD	6.911	7.121	8.011	8.739
EMUL/CMD (H)	15.140	15.386	15.864	16.116
EMUL/CMD	15.359	15.795	16.293	16.604
Processor Util.				
TOTAL (H)	87.33	176.49	269.55	522.88
TOTAL	87.00	177.00	270.00	522.00
UTIL/PROC (H)	87.33	88.25	89.85	87.15
UTIL/PROC	87.00	88.50	90.00	87.00
TVR(H)	1.48	1.48	1.53	1.58
TVR	1.45	1.45	1.49	1.53
Storage				
WKSET (V)	70	66	63	58
PGBLPGS	25796	54393	83039	104K
PGBLPGS/USER	23.5	25.2	26.6	18.2
FREEPGS	2803	5336	7627	14349
FREE UTIL	0.94	0.96	0.97	0.96
SHRPGS	801	886	995	1186
Paging				
READS/SEC	170	224	208	332
WRITES/SEC	90	112	91	186
PAGE/CMD	6.655	4.350	2.691	2.515
XSTOR IN/SEC	127	295	479	1053
XSTOR OUT/SEC	233	437	597	1298
XSTOR/CMD	9.215	9.477	9.686	11.414
FAST CLR/CMD	5.222	5.425	5.509	5.651

RELEASE RUN ID	ESA 1.1 Y14R1101	ESA 1.1 Y24R2161	ESA 1.1 Y34R3122	ESA 1.1 Y64R5865
Environment				
REAL STORAGE	**128M**	**256M**	**384M**	**512M**
EXP. STORAGE	**256M**	**512M**	**1024M**	**2048M**
USERS	**1100**	**2160**	**3118**	**5860**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**0**	**0**	**2**	**2**
PROCESSORS	**1**	**2**	**3**	**6**
I/O				
VIO RATE	324	654	948	1751
VIO/CMD	8.294	8.467	8.534	8.501
MDC READS	220	435	637	1223
MDC WRITES	102	209	295	550
MDC MODS	84	171	244	463
MDC HIT RATIO	0.92	0.92	0.92	0.93
PRIVOPs				
PRIVOP/CMD	17.624	17.490	22.279	20.075
DIAG/CMD	24.746	24.681	24.315	23.603
DIAG 08/CMD	0.742	0.738	0.765	0.719
DIAG 10/CMD	0.000	0.013	0.009	0.015
DIAG 58/CMD	1.254	1.243	1.251	1.219
DIAG 98/CMD	0.384	0.401	0.405	0.291
DIAG A4/CMD	3.916	3.975	3.997	4.083
DIAG A8/CMD	1.843	1.955	1.989	1.893
DIAG 214/CMD	12.184	12.209	12.305	12.371
SIE/CMD	53.780	53.392	59.348	54.688
SIE INTCPT/CMD	36.033	35.239	40.357	35.547
FREE TOTL/CMD	80.735	80.334	84.498	84.517
VTAM Machines				
WKSET (V)	818	1611	1319	1791
TOT CPU/CMD (V)	3.0717	3.0209	3.9657	3.8211
CP CPU/CMD (V)	1.3368	1.3319	1.9804	1.8823
VIRT CPU/CMD (V)	1.7349	1.6890	1.9854	1.9388
DIAG 98/CMD (V)	0.393	0.405	0.411	0.295
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Table 31. VM/ESA 1.1 on Selected 9021 Processors.

9121 / Minidisk

The following is a description of the environments used to test VM/ESA 1.1 processor capacity on the 9121-480 processor.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

- PROCESSOR:

9121-320:	9121-480	CPU 1 only
9121-480:	9121-480	CPUs 0-1

- STORAGE:
  - RSTOR: 192M
  - XSTOR: 64M All reserved for MDC

- DASD:

	PACK NAME	TYPE
- SYSTEM:	PSYS02	3380-A
	PSPT01	3380-A
	WKLD01	3380-A
	WKLD02	3380-A

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
9121-320:						
3380-A	11 - 3880-2	16	4	8	10	0
3380-A	2 - 3880-J23	0	0	0	4	0
9121-480:						
3380-A	11 - 3880-2	16	4	8	20	0
3380-A	2 - 3880-J23	0	0	0	4	0

(NUMBER OF PACKS: PAGE, SPOOL, TDISK, USER, SERVER)

- TAPE: MONITOR 3480

- COMMUNICATIONS:

	CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
9121-320:	3745-410	1	44	56Kb
9121-480:	3745-410	2	44	56Kb

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2M
- USER CMS MODE: XA
- USER RELSHARE: 100

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
9121-320:				
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
9121-480:				
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON, RESERVE 1050

4) MEASUREMENT DISCUSSION

Two 9121 measurements were obtained. The hardware differences between the two runs were:

- Only CPU 1 online for first run (making it equivalent to 9121-320)
- Additional DASD volumes for the second run to accommodate the additional users
- Extra 3745 used for the second run to accommodate the additional users

Both runs were equivalent in software setup except the second run had the VSCS server′s working set reserved in storage with the SET RESERVE command. This was done because running the FS7B0R workload at a 90% processor utilization on the 9121-480 was moderately storage constrained and it was desirable to ensure that the VSCS servers′ pages remained in storage to improve the users′ response times.

The number of users was adjusted so that both runs had approximately the same processor utilization. The utilization per processor (UTIL/PROC (H)) was 87.75 and 87.30%, respectively. Both runs were completed with the same amount of storage and there was an increase of approximately one page I/O per command (PAGE/CMD). The ITR ratio (ITRR (H)) shows that the ITR of the 9121-480 run is almost double the rate of the 9121-320 run. This result is consistent with what has been observed for the HT4 and PD3 hardware capacity workloads that are used to do processor evaluation for the VM CMS intensive environment.

PROCESSOR MODEL RELEASE RUN ID	9121-320 ESA 1.1 L14R0910	9121-480 ESA 1.1 L24R1770
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**910**	**1770**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**1**	**2**
Response Time		
TRIV INT	0.068	0.065
NONTRIV INT	0.453	0.410
TOT INT	0.319	0.273
TOT INT ADJ	0.329	0.301
AVG FIRST (T)	0.280	0.355
AVG LAST (T)	0.490	0.535
Throughput		
AVG THINK (T)	25.82	25.74
ETR	33.39	69.14
ETR (T)	32.41	62.66
ETR RATIO	1.030	1.103
ITR (H)	36.94	71.78
ITR	38.09	39.62
EMUL ITR	58.55	61.15
ITRR (H)	1.000	1.943
ITRR	1.000	1.040
Proc. Usage		
PBT/CMD (H)	27.070	27.863
PBT/CMD	27.148	27.927
CP/CMD (H)	9.784	10.176
CP/CMD	9.564	9.894
EMUL/CMD (H)	17.279	17.681
EMUL/CMD	17.585	18.033
Processor Util.		
TOTAL (H)	87.75	174.59
TOTAL	88.00	175.00
UTIL/PROC (H)	87.75	87.30
UTIL/PROC	88.00	87.50
TVR(H)	1.57	1.58
TVR	1.54	1.55
Storage		
WKSET (V)	81	74
PGBLPGS	43242	39871
PGBLPGS/USER	47.5	22.5
FREEPGS	2373	4505
FREE UTIL	0.93	0.97
SHRPGS	784	891
Paging		
READS/SEC	252	507
WRITES/SEC	159	340
PAGE/CMD	12.679	13.517
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.399	5.458

PROCESSOR MODEL RELEASE RUN ID	9121-320 ESA 1.1 L14R0910	9121-480 ESA 1.1 L24R1770
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**910**	**1770**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**1**	**2**
I/O		
VIO RATE	279	541
VIO/CMD	8.607	8.634
MDC READS	187	359
MDC WRITES	88	169
MDC MODS	72	139
MDC HIT RATIO	0.92	0.92
PRIVOPs		
PRIVOP/CMD	24.160	21.869
DIAG/CMD	25.339	25.202
DIAG 08/CMD	0.771	0.766
DIAG 10/CMD	0.000	0.016
DIAG 58/CMD	1.234	1.245
DIAG 98/CMD	0.432	0.463
DIAG A4/CMD	4.041	4.038
DIAG A8/CMD	2.005	2.043
DIAG 214/CMD	12.340	12.320
SIE/CMD	61.485	58.536
SIE INTCPT/CMD	44.884	40.975
FREE TOTL/CMD	86.720	84.309
VTAM Machines		
WKSET (V)	1055	1217
TOT CPU/CMD (V)	4.8675	4.4861
CP CPU/CMD (V)	2.4337	2.2253
VIRT CPU/CMD (V)	2.4337	2.2608
DIAG 98/CMD (V)	0.461	0.466
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 32. 9121 Processor Capacity

Storage Constrained Runs

9021-720 / 35% SFS

The following is a description of the environment used to obtain the various storage constrained runs on a 9021-720 running VM/ESA 1.1. The FS7B35R CMS intensive workload with the same hardware configuration was used in all cases with the exception of the amount of real and expanded storage available to the system.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:   9021-720
- STORAGE:
   - RSTOR:   256M-512M (See table)
   - XSTOR:     512M-2G (See table)

- DASD:
             PACK NAME   TYPE
   - SYSTEM: PSYS02      3380-A
             PSPT01      3380-A
             WKLD01      3380-A
             WKLD02      3380-A

    TYPE OF    NUMBER/TYPE OF            NUMBER OF PACKS
    DASD       CONTROL UNIT     PAGE   SPOOL   TDISK   USER    SERVER
    3380-A      12 - 3880-2      20     16      12      0        0
    3380-A       4 - 3990-2       0      0       0      0       18

- TAPE:       MONITOR  3480

- COMMUNICATIONS:
                  CONTROLLER  NUMBER  LINES/CONTROLLER  LINESPEED
                     3745       3           44             56Kb
```

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:   100

- SERVER MACHINES:
                                VM SIZE/
 SERVER MACHINE     TYPE       CMS MODE   RELSHARE OTHER OPTIONS
 VTAMXAA            VTAM        64M/XA     10000   QUICKDSP ON
 VSCSXA2            VSCS        64M/XA     10000   QUICKDSP ON
 VSCSXA3            VSCS        64M/XA     10000   QUICKDSP ON
 SERVE2             SFS         32M/XA      1500   QUICKDSP ON
 SERVE4             SFS         32M/XA      1500   QUICKDSP ON
 SERVE7             SFS         32M/XA      1500   QUICKDSP ON
 SERVE8             SFS         32M/XA      1500   QUICKDSP ON
 CRRSERVA           CRR         17M/XA       100   NONE
```

4) MEASUREMENT DISCUSSION

A set of six runs was completed on the 9021-720 processor to determine how well VM/ESA 1.1 performed in storage constrained environments and what the minimum storage requirements would be to run the FS7B35R workload with 4800 users and still achieve “acceptable” performance. For the purposes of this discussion, performance will be considered acceptable when the external response time is less than one second. All runs were completed on an early level of code that is similar in performance to the GA level code. All runs were completed with the same hardware and software configuration except for the real and expanded storage sizes.

As expected, the performance of VM improved with the amount of storage available to the system. External response times decreased and throughput (ITR (H) and ETR(T)) increased as the available storage increased.

The reader may note that the RTM based ITR did not follow this trend. This is due to a limitation in the way CP determines when a transaction begins and ends. What was happening with the first three runs in the table is that CP was double counting transactions when users were put into page wait. It determines that a given virtual machine is run twice and counts this as two transactions when there actually is only one transaction that is put into page wait and re-run.

As paging rates and response times increased, more and more users were being put in the eligible list (AVG ELIST SIZE (V)). In fact, eligible list formation is an indicator of a system that is storage constrained. The average working set size, a function of both the actual working set size and the demand on storage decreased as the storage size decreased. The first three runs (the most constrained runs) in the table show an increase in TVRs because of the increased CP overhead to support paging. The reader may notice that even in the most constrained environment, the minidisk cache hit ratio was above 80%. In all runs, the minimum minidisk cache size was set to 64M which appears ample enough for the FS7B35R workload to achieve a good hit ratio.

The reader should also refer to the graphs and related information in “Storage Constrained Runs” on page 34 for the following discussion.

It is important to note that there was an abrupt transition between acceptable and unacceptable performance. Run Y64F480E conducted with 320M real and 1024M expanded storage exhibited an acceptable external response time (AVG LAST (T)) of 0.8 seconds, while taking away just 128M of expanded storage for run Y64D480D had an external response time of 6.8 seconds. The following table shows the percentage increase for the following statistics when moving from the first run to the second:

Run Id: Real/Exp.	Y64F4809 - Y64F480E 384M/1G - 320M/1G		Y64F480E - Y64F480D 320M/1G - 320M/896M		Y64F480A - Y64F4808 256M/768M - 256M/512M	
	Delta	Pct.	Delta	Pct.	Delta	Pct.
AVG LAST (T)	0.104	14.8%	6.007	747.0%	53.479	547.5%
PBT/CMD	0.410	1.3%	4.381	13.8%	44.094	105.4%
PAGE/CMD	0.959	10.5%	3.047	30.2%	22.447	161%
XSTOR/CMD	2.832	22.1%	21.264	136.0%	30.353	47.6%

Table 33. Comparing Specific Storage Constrained Runs

From the first column in this table, it is clear that the removal of real storage increased the paging rate but PBT/CMD only rose by 1.3% and didn′t entail much extra CP overhead to manage paging. Response times did increase but were still sub-second. The second column shows that removing 128M of expanded storage caused the expanded storage paging rate to greatly increase. It appears that there was no longer enough expanded storage to accommodate the number of active users on the system. PBT/CMD increased in this case by over 13% as CP was spending much more time paging and response times increased 747%. Thus, when running in a storage constrained environment such as in run Y64F480E, a small reduction in the amount of storage and/or an increase in the storage requirements of the workload can push the performance of VM past the acceptable point and dramatically degrade performance.

The last column of the preceding table shows the performance difference between the two most storage constrained runs. External response time for the most constrained run was over 54 seconds longer than the next most constrained run. This huge jump in response time was due to the huge increase in the DASD page I/Os per command (PAGE/CMD) which resulted in a much higher processor busy time per command (PBT/CMD), mostly in CP to manage paging.

RELEASE RUN ID	ESA 1.1 Y64F4808	ESA 1.1 Y64F480A	ESA 1.1 Y64F480D	ESA 1.1 Y64F480E	ESA 1.1 Y64F4809	ESA 1.1 Y64F480F
Environment						
REAL STORAGE	**256M**	**256M**	**320M**	**320M**	**384M**	**512M**
EXP. STORAGE	**512M**	**768M**	**896M**	**1024M**	**1024M**	**2048M**
USERS	**4800**	**4800**	**4800**	**4800**	**4800**	**4800**
VTAMs	**1**	**1**	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**	**6**	**6**
Response Time						
TRIV INT	5.075	0.084	0.057	0.076	0.070	0.045
NONTRIV INT	26.320	1.778	1.242	0.542	0.485	0.324
TOT INT	13.748	0.714	0.553	0.342	0.310	0.208
TOT INT ADJ	31.546	1.505	1.013	0.399	0.355	0.232
AVG FIRST (T)	48.388	7.306	4.740	0.515	0.460	0.330
AVG LAST (T)	63.246	9.767	6.811	0.804	0.700	0.502
Throughput						
AVG THINK (T)	24.82	25.55	25.52	25.23	25.27	25.62
ETR	119.65	278.27	259.17	196.70	193.39	189.73
ETR (T)	52.14	131.99	141.53	168.70	168.99	169.84
ETR RATIO	2.295	2.108	1.831	1.166	1.144	1.117
ITR (H)	69.55	143.41	166.25	189.57	192.06	na
ITR	26.73	50.46	50.78	36.83	36.61	37.66
EMUL ITR	96.14	94.47	87.47	60.42	59.43	59.00
ITRR (H)	1.000	2.062	2.390	2.726	2.761	na
ITRR	1.000	1.888	1.899	1.378	1.370	1.409
Proc. Usage						
PBT/CMD (H)	86.266	41.837	36.089	31.651	31.240	na
PBT/CMD	85.916	41.822	36.035	31.654	31.244	29.675
CP/CMD (H)	64.426	21.139	16.139	12.878	12.508	na
CP/CMD	62.135	19.471	15.121	12.330	11.953	10.716
EMUL/CMD (H)	21.835	20.693	19.945	18.768	18.726	na
EMUL/CMD	23.780	22.350	20.914	19.325	19.291	18.959
Processor Util.						
TOTAL (H)	449.83	552.20	510.77	533.95	527.94	na
TOTAL	448.00	552.00	510.00	534.00	528.00	504.00
UTIL/PROC (H)	74.97	92.03	85.13	88.99	87.99	na
UTIL/PROC	74.67	92.00	85.00	89.00	88.00	84.00
TVR(H)	3.95	2.02	1.81	1.69	1.67	na
TVR	3.61	1.87	1.72	1.64	1.62	1.57
Storage						
WKSET (V)	24	31	45	66	70	63
PGBLPGS	40663	42003	58639	59910	75438	108K
PGBLPGS/USER	8.5	8.8	12.2	12.5	15.7	23.0
FREEPGS	14276	13249	13061	12589	12579	12407
FREE UTIL	0.93	0.95	0.94	0.96	0.96	0.97
SHRPGS	1147	1306	1355	1424	1370	1333
Paging						
AVG ELIST SIZE (V)	204	37	2	0	0	0
READS/SEC	1048	1011	1029	954	876	289
WRITES/SEC	849	828	831	749	668	161
PAGE/CMD	36.380	13.933	13.142	10.095	9.136	2.650
XSTOR IN/SEC	2076	3745	2220	891	695	1063
XSTOR OUT/SEC	2830	4667	3006	1751	1473	1282
XSTOR/CMD	94.085	63.732	36.925	15.661	12.829	13.807
FAST CLR/CMD	5.715	5.675	5.737	5.578	5.598	5.711

RELEASE RUN ID	ESA 1.1 Y64F4808	ESA 1.1 Y64F480A	ESA 1.1 Y64F480D	ESA 1.1 Y64F480E	ESA 1.1 Y64F4809	ESA 1.1 Y64F480F
Environment						
REAL STORAGE	**256M**	**256M**	**320M**	**320M**	**384M**	**512M**
EXP. STORAGE	**512M**	**768M**	**896M**	**1024M**	**1024M**	**2048M**
USERS	**4800**	**4800**	**4800**	**4800**	**4800**	**4800**
VTAMs	**1**	**1**	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**	**6**	**6**
I/O						
VIO RATE	353	922	1013	1171	1174	1175
VIO/CMD	6.770	6.985	7.158	6.941	6.947	6.918
MDC READS	284	735	795	931	934	1042
MDC WRITES	112	249	268	301	313	312
MDC MODS	50	89	98	121	137	242
MDC HIT RATIO	0.82	0.82	0.82	0.83	0.84	0.93
PRIVOPs						
PRIVOP/CMD	28.420	26.458	27.352	27.334	27.449	27.693
DIAG/CMD	20.828	20.921	21.402	20.806	20.811	20.811
DIAG 08/CMD	0.690	0.742	0.763	0.735	0.740	0.730
DIAG 10/CMD	0.019	0.015	0.014	0.012	0.012	0.012
DIAG 58/CMD	1.227	1.235	1.272	1.239	1.231	1.236
DIAG 98/CMD	0.403	0.341	0.339	0.308	0.308	0.306
DIAG A4/CMD	2.589	2.727	2.784	2.727	2.734	2.732
DIAG A8/CMD	1.688	1.735	1.795	1.695	1.698	1.666
DIAG 214/CMD	11.123	11.478	11.757	11.512	11.497	11.564
SIE/CMD	104.633	93.098	79.588	72.841	66.653	72.350
SIE INTCPT/CMD	47.085	45.618	46.161	48.803	45.991	49.921
FREE TOTL/CMD	132.383	100.857	101.294	97.121	96.950	96.466
VTAM Machines						
WKSET (V)	1540	1781	1719	1802	1823	1555
TOT CPU/CMD (V)	8.9659	5.1267	4.5588	3.8871	3.8594	3.8707
CP CPU/CMD (V)	4.9089	2.3823	2.1460	1.9090	1.8903	1.8896
VIRT CPU/CMD (V)	4.0570	2.7443	2.4128	1.9781	1.9692	1.9811
DIAG 98/CMD (V)	0.408	0.344	0.340	0.312	0.308	0.311
SFS Servers						
WKSET (V)	521	599	891	980	1053	992
TOT CPU/CMD (V)	5.5691	4.8741	4.5117	4.1175	4.1389	3.8119
CP CPU/CMD (V)	3.0028	2.3613	2.1617	2.0045	2.0119	1.8733
VIRT CPU/CMD (V)	2.5662	2.5128	2.3500	2.1131	2.1270	1.9386
FP REQ/CMD(Q)	1.359	1.355	1.380	1.337	1.345	1.348
IO/CMD (Q)	2.421	2.256	2.203	2.014	2.018	2.003
IO TIME/CMD (Q)	0.707	0.174	0.174	0.071	0.068	0.037
SFS TIME/CMD (Q)	1.694	0.407	0.392	0.138	0.129	0.085

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

Table 34. VM/ESA 1.1 on 9021-720 with Various Amounts of Storage

9. New Functional Enhancements

VM Data Spaces: Exploitation by Shared File System

In Release 1.1, SFS exploits data spaces through read only access to DIRCONTROL directories. The SFS server, on first ACCESS of a DIRCONTROL directory that has been identified as eligible for data spaces, builds the FSTs within a data space and maps the file data to the data space. Subsequent data reads are done via direct reference to the data space.

The measurements in the following sections show SFS exploitation of VM Data Spaces by comparing the performance of SFS with data spaces to other environments on the 9021-720, 9121-480, and the 9221-170. The first three subsections have the read/write data in SFS while the last subsection has the read/write data on minidisks.

9021-720 / XC Mode Users

1) WORKLOAD: FS7B35R and FS7BMAXR

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-720
- STORAGE:
   - RSTOR:   512M
   - XSTOR:   2G

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              PSPT01     3380-D
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-K	4 - 3990-02	0	0	0	0	16
3380-A	10 - 3880-03	20	8	12	0	0

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	22	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XC
- USER RELSHARE:    100
```

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
SERVE2 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
SERVE4 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
SERVE7 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
SERVE8 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
SERVE6 (R/O)	SFS	32M/XC	1500	QUICKDSP ON
CRRSERVA	CRR	32M/XC	100	NONE

4) MEASUREMENT DISCUSSION

For all measurements shown in this section, the user virtual machines were run in XC mode. Virtual machines running in XC mode directly reference a single shared copy of the FSTs in the data space. The file data is moved from the data space by going into access register mode and using the MVCL instruction.

All environments measured have the read/write data in SFS directories but vary the location of the read-only data. The following describes the location of the R/O data for each measurement:

Y64F480L R/O data on minidisks. One of the four minidisks has its FSTs in a shared segment. Minidisk caching is in effect. This case is intended to represent a typical usage of read-only minidisks.

Y64M480J R/O data in SFS DIRCONTROL directories in data spaces. These directories reside in a separate file pool, as recommended in the *VM/ESA 1.1 CMS Planning and Administration Guide*.

Y64M480K R/O data on minidisks; all four minidisks have their FSTs in shared segments. Minidisk caching is in effect. This case represents best case minidisk performance.

Y64M480O R/O data in SFS FILECONTROL directories which reside in a separate file pool. Minidisk caching is in effect.

When comparing the SFS data in data spaces environment to minidisks with all FSTs in shared segments, there was an increase in processor busy time (PBT/CMD (H)) of 1.6%. This is in part due to the processing required to “hook” the user machine to the data space and Coordinated Resource Recovery (CRR) processing in accessing multiple file pools.

Relative to the case where only one minidisk has its FSTs in a shared segment, processor busy time and external response time (AVG LAST (T)) were about the same. The benefit of shared FSTs is evident as there was a 23% reduction in paging (XSTOR/CMD). This tends to counterbalance the additional processing cited previously.

Relative to the case where the R/O data is in SFS but not in data spaces, there was a decrease in processor busy time (PBT/CMD (H)) of 6.1%. This is because the overhead of communicating to the server and the SFS server processing associated with normal SFS usage is essentially eliminated. Additionally, data spaces provided a 21% reduction in paging (XSTOR/CMD) because of the shared FSTs.

R/O Data # Shared FSTs RELEASE RUN ID	Minidisk 1 OF 4 ESA 1.1 Y64F480L	DIRC w/ D.S. 4 OF 4 ESA 1.1 Y64M480J	Minidisk 4 OF 4 ESA 1.1 Y64M480K	FILECONTROL 0 OF 4 ESA 1.1 Y64M480O
Environment				
REAL STORAGE	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**
USERS	**4800**	**4800**	**4800**	**4800**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**
Response Time				
TRIV INT	0.049	0.048	0.046	0.054
NONTRIV INT	0.346	0.343	0.335	0.392
TOT INT	0.222	0.221	0.216	0.248
TOT INT ADJ	0.249	0.248	0.240	0.286
AVG FIRST (T)	0.357	0.345	0.333	0.403
AVG LAST (T)	0.525	0.517	0.495	0.602
Throughput				
AVG THINK (T)	25.61	25.52	25.69	25.64
ETR	190.35	190.32	188.52	194.83
ETR (T)	169.39	169.58	169.34	168.99
ETR RATIO	1.124	1.122	1.113	1.153
ITR (H)	199.92	199.48	202.59	187.21
ITR	37.48	37.38	37.56	35.97
EMUL ITR	58.88	58.24	58.82	57.00
ITRR (H)	1.000	0.998	1.013	0.936
ITRR	1.000	0.997	1.002	0.960
Proc. Usage				
PBT/CMD (H)	30.012	30.078	29.616	32.049
PBT/CMD	29.991	30.015	29.644	32.074
CP/CMD (H)	11.445	11.215	11.166	12.329
CP/CMD	10.922	10.732	10.689	11.835
EMUL/CMD (H)	18.561	18.857	18.444	19.714
EMUL/CMD	19.069	19.282	18.956	20.238
Processor Util.				
TOTAL (H)	508.36	510.07	501.52	541.59
TOTAL	508.00	509.00	502.00	542.00
UTIL/PROC (H)	84.73	85.01	83.59	90.27
UTIL/PROC	84.67	84.83	83.67	90.33
TVR(H)	1.62	1.59	1.61	1.63
TVR	1.57	1.56	1.56	1.58
Storage				
WKSET (V)	64	54	52	60
PGBLPGS	108K	107K	107K	108K
PGBLPGS/USER	23.0	22.8	22.8	23.0
FREEPGS	12229	12372	12413	12188
FREE UTIL	0.96	0.95	0.95	0.95
SHRPGS	1354	1338	1386	1388
Paging				
READS/SEC	298	214	232	287
WRITES/SEC	163	91	105	159
PAGE/CMD	2.722	1.799	1.990	2.639
XSTOR IN/SEC	1086	848	795	1039
XSTOR OUT/SEC	1304	989	957	1267
XSTOR/CMD	14.110	10.832	10.346	13.646
FAST CLR/CMD	5.632	5.743	5.651	5.793

R/O Data # Shared FSTs RELEASE RUN ID	Minidisk 1 OF 4 ESA 1.1 Y64F480L	DIRC w/ D.S. 4 OF 4 ESA 1.1 Y64M480J	Minidisk 4 OF 4 ESA 1.1 Y64M480K	FILECONTROL 0 OF 4 ESA 1.1 Y64M480O
Environment				
REAL STORAGE	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**
USERS	**4800**	**4800**	**4800**	**4800**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**	**6**
I/O				
VIO RATE	1174	1095	1165	1087
VIO/CMD	6.931	6.457	6.880	6.433
MDC READS	1041	925	1044	1008
MDC WRITES	310	306	298	311
MDC MODS	241	245	242	241
MDC HIT RATIO	0.93	0.93	0.94	0.93
PRIVOPs				
PRIVOP/CMD	28.764	29.018	28.944	31.649
DIAG/CMD	21.276	20.764	21.339	20.814
DIAG 08/CMD	0.738	0.737	0.732	0.746
DIAG 10/CMD	0.012	0.012	0.012	0.012
DIAG 58/CMD	1.240	1.238	1.240	1.243
DIAG 98/CMD	0.313	0.313	0.319	0.314
DIAG A4/CMD	2.745	2.223	2.716	2.213
DIAG A8/CMD	1.659	1.698	1.636	1.687
DIAG 214/CMD	11.601	11.546	11.627	11.640
SIE/CMD	72.545	66.421	66.517	72.716
SIE INTCPT/CMD	50.056	47.159	47.227	50.902
FREE TOTL/CMD	96.726	96.613	96.752	103.015
VTAM Machines				
WKSET (V)	1446	1436	1475	1438
TOT CPU/CMD (V)	3.9795	3.9716	3.9651	3.9188
CP CPU/CMD (V)	1.9865	1.9776	1.9876	1.9364
VIRT CPU/CMD (V)	1.9930	1.9940	1.9775	1.9824
DIAG 98/CMD (V)	0.316	0.316	0.320	0.319
SFS Servers				
WKSET (V)	935	1904	1063	1993
TOT CPU/CMD (V)	3.6943	3.7948	3.7141	5.0563
CP CPU/CMD (V)	1.7865	1.8303	1.7977	2.3539
VIRT CPU/CMD (V)	1.9078	1.9645	1.9164	2.7024
FP REQ/CMD(Q)	1.339	1.401	1.352	2.186
IO/CMD (Q)	1.954	1.990	1.986	2.402
IO TIME/CMD (Q)	0.036	0.035	0.034	0.039
SFS TIME/CMD (Q)	0.087	0.083	0.086	0.097

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

Table 35. SFS Usage of VM Data Spaces - XC Mode (9021-720)

9021-720 / 370 Mode Users

1) WORKLOAD: FS7B35R and FS7BMAXR

2) HARDWARE CONFIGURATION

- PROCESSOR: 9021-720
- STORAGE:
  - RSTOR: 512M
  - XSTOR: 2G

- DASD:

	PACK NAME	TYPE
SYSTEM:	PSYS02	3380-A
	PSPT01	3380-D
	WKLD01	3380-D
	WKLD02	3380-D

DASD	CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-K	4 - 3990-02	0	0	0	0	16
3380-A	10 - 3880-03	20	8	12	0	0

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	22	56Kb

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2176K
- USER CMS MODE: 370
- USER RELSHARE: 100

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
SERVE2 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
SERVE4 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
SERVE7 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
SERVE8 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
SERVE6 (R/O)	SFS	32M/XC	1500	QUICKDSP ON
CRRSERVA	CRR	32M/XC	100	NONE

4) MEASUREMENT DISCUSSION

For the measurements shown in this section, the user virtual machines were run in 370 mode. Virtual machines in 370 or XA mode gain a performance advantage as the already built FSTs are copied into the virtual machine using the CP DIAGNOSE X¢248¢ instruction. The file data blocks are also moved into the virtual machine using the DIAGNOSE instruction, thus eliminating server communication and processing associated with normal SFS usage. Relative to XC mode, the main significant disadvantage is that it is no longer practical to share the FSTs among users. This disadvantage is more significant in environments having limited or no expanded storage for paging.

Both environments measured have the read/write data in SFS directories but vary the location of the read-only data:

Y64F480M R/O data on minidisks. One of the four minidisks has its FSTs in a shared segment. Minidisk caching is in effect. This case is intended to represent a typical usage of read-only minidisks.

Y64M480P R/O data in SFS DIRCONTROL directories in data spaces. These directories reside in a separate file pool.

The results showed equivalent external response time (AVG LAST (T)) but a slightly lower internal throughput rate (ITR (H) = -0.6%) for the data space environment.

When comparing the 370 mode data space run to the XC mode data space run in "9021-720 / XC Mode Users" on page 161, there was a decrease in processor busy time per command (PBT/CMD (H)) of 1.7%. This is primarily due to the fact that somewhat less processing is required to run in 370 mode than in XC mode. For more information see section "Software Mode Comparisons" on page 119. The increase of 22% in paging in the 370 mode case was due to the fact that the FSTs are not shared. The absence of shared FSTs has only a minimal adverse effect on performance in this environment because most paging can be done to expanded storage.

R/O Data # Shared FSTs RELEASE RUN ID	Minidisk 1 OF 4 ESA 1.1 Y64F480M	DIRCONTROL w/ D.S. 4 OF 4 ESA 1.1 Y64M480P
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**4800**	**4800**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
Response Time		
TRIV INT	0.052	0.051
NONTRIV INT	0.359	0.347
TOT INT	0.238	0.231
TOT INT ADJ	0.259	0.252
AVG FIRST (T)	0.333	0.335
AVG LAST (T)	0.497	0.495
Throughput		
AVG THINK (T)	25.66	25.62
ETR	185.18	184.66
ETR (T)	169.93	169.57
ETR RATIO	1.090	1.089
ITR (H)	204.15	202.85
ITR	37.13	36.93
EMUL ITR	59.33	58.91
ITRR (H)	1.000	0.994
ITRR	1.000	0.995
Proc. Usage		
PBT/CMD (H)	29.391	29.579
PBT/CMD	29.424	29.486
CP/CMD (H)	11.537	11.489
CP/CMD	11.005	10.969
EMUL/CMD (H)	17.847	18.084
EMUL/CMD	18.420	18.517
Processor Util.		
TOTAL (H)	499.42	501.58
TOTAL	500.00	500.00
UTIL/PROC (H)	83.24	83.60
UTIL/PROC	83.33	83.33
TVR(H)	1.65	1.64
TVR	1.60	1.59
Storage		
WKSET (V)	64	60
PGBLPGS	108K	108K
PGBLPGS/USER	23.0	23.0
FREEPGS	12189	12410
FREE UTIL	0.97	0.95
SHRPGS	1362	1358
Paging		
READS/SEC	302	300
WRITES/SEC	173	164
PAGE/CMD	2.795	2.736
XSTOR IN/SEC	1086	1006
XSTOR OUT/SEC	1313	1245
XSTOR/CMD	14.118	13.275
FAST CLR/CMD	5.467	5.708

R/O Data # Shared FSTs RELEASE RUN ID	Minidisk 1 OF 4 ESA 1.1 Y64F480M	DIRCONTROL w/ D.S. 4 OF 4 ESA 1.1 Y64M480P
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**4800**	**4800**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
I/O		
VIO RATE	1176	1087
VIO/CMD	6.921	6.410
MDC READS	1052	905
MDC WRITES	312	307
MDC MODS	242	238
MDC HIT RATIO	0.93	0.92
PRIVOPs		
PRIVOP/CMD	28.804	28.918
DIAG/CMD	21.496	21.367
DIAG 08/CMD	0.724	0.725
DIAG 10/CMD	0.012	0.012
DIAG 58/CMD	1.242	1.238
DIAG 98/CMD	0.312	0.313
DIAG A4/CMD	2.748	2.211
DIAG A8/CMD	1.654	1.681
DIAG 214/CMD	11.240	11.140
SIE/CMD	72.313	72.464
SIE INTCPT/CMD	51.343	51.450
FREE TOTL/CMD	102.444	102.658
VTAM Machines		
WKSET (V)	1435	1444
TOT CPU/CMD (V)	3.9662	4.0035
CP CPU/CMD (V)	1.9831	1.9952
VIRT CPU/CMD (V)	1.9831	2.0083
DIAG 98/CMD (V)	0.313	0.316
SFS Servers		
WKSET (V)	966	1774
TOT CPU/CMD (V)	3.7639	3.7938
CP CPU/CMD (V)	1.8268	1.8347
VIRT CPU/CMD (V)	1.9371	1.9592
FP REQ/CMD(Q)	1.340	1.394
IO/CMD (Q)	1.974	1.948
IO TIME/CMD (Q)	0.037	0.036
SFS TIME/CMD (Q)	0.087	0.084

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

9121-480 / XC Mode Users

1) WORKLOAD: FS7B35R and FS7BMAXR

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9121-480
- STORAGE:
   - RSTOR:       192M
   - XSTOR:        64M     All reserved for MDC

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380
              PSPT01     3380
              WKLD01     3380
              WKLD02     3380
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER/SERVER
		NUMBER OF PACKS			
3380-A	11 - 3880-2	16	4	8	20
3380-A	2 - 3880-J23	0	0	0	4

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	2	44	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XC
- USER RELSHARE:    100
```

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON, RESERVE 850
RWSERV1 (R/W)	SFS	32M/XA	1500	QUICKDSP ON, RESERVE 1300
RWSERV2 (R/W)	SFS	32M/XA	1500	QUICKDSP ON, RESERVE 1300
ROSERV1 (R/O)	SFS	32M/XA	1500	QUICKDSP ON
CRRSERV1	CRR	17M/XA	100	NONE

Note: SET RESERVE was specified for VSCS and the two SFS production servers to avoid a serial page bottleneck problem within the servers. The number of pages reserved was set equal to that virtual machine's working set size. See section "Set Reserve Option" on page 232 for details.

4) MEASUREMENT DISCUSSION

All users are run in XC mode. All environments measured have the read/write data in SFS directories but vary the location of the read-only data. Paging to expanded storage is not performed in these environments. The following describes the location of the R/O data for each measurement:

L24F1483 R/O data on minidisks. One of the four minidisks has its FSTs in a shared segment. Minidisk caching is in effect. This case is intended to represent a typical usage of read-only minidisks.

L24F1481 R/O data in SFS DIRCONTROL directories in data spaces. These directories reside in a separate file pool.

L24F1482 R/O data on minidisks; one of the four minidisks has its FSTs in a shared segment. Minidisk caching is *NOT* in effect for the R/O minidisks and they are not behind cached control units.

The 9121-480 environment, which does not have expanded storage for paging, benefits more than the 9021-720 case from having shared FSTs. When comparing the data spaces environment to minidisks with MDC, there was a decrease in paging (PAGE/CMD) of 22% and a 14% improvement in external response time (AVG LAST (T)). Based on the 9021-720 measurements, it can also be expected that 370 mode data space usage will not do as well in this environment. Not having shared FSTs had a more adverse performance impact in this environment where there is no expanded storage for paging.

Use of data spaces showed a 29% improvement in external response time relative to the case where the R/O data is currently on minidisks without minidisk or control unit caching. Data spaces provide an additional advantage in this environment because the file data becomes cached in real memory instead of having to be read from DASD.

R/O Data # Shared FSTs MINIDISK CACHING CONTROL UNIT CACHING RELEASE RUN ID	Minidisk 1 OF 4 YES YES ESA 1.1 L24F1483	DIRCONTROL w/ D.S. 4 OF 4 YES YES ESA 1.1 L24F1481	Minidisk 1 OF 4 NO NO ESA 1.1 L24F1482
Environment			
REAL STORAGE	**192M**	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**	**64M**
USERS	**1480**	**1480**	**1480**
VTAMs	**1**	**1**	**1**
VSCSs	**1**	**1**	**1**
PROCESSORS	**2**	**2**	**2**
Response Time			
TRIV INT	0.067	0.059	0.069
NONTRIV INT	0.630	0.595	0.712
TOT INT	0.410	0.392	0.451
TOT INT ADJ	0.454	0.419	0.516
AVG FIRST (T)	0.420	0.330	0.515
AVG LAST (T)	0.710	0.610	0.855
Throughput			
AVG THINK (T)	25.59	25.49	25.53
ETR	58.34	56.53	59.84
ETR (T)	52.64	52.93	52.35
ETR RATIO	1.108	1.068	1.143
ITR (H)	60.67	60.76	60.93
ITR	33.66	32.48	34.85
EMUL ITR	53.51	50.88	55.46
ITRR (H)	1.000	1.002	1.004
ITRR	1.000	0.965	1.035
Proc. Usage			
PBT/CMD (H)	32.968	32.917	32.827
PBT/CMD	32.867	32.875	32.859
CP/CMD (H)	12.599	12.270	12.581
CP/CMD	12.159	11.903	12.227
EMUL/CMD (H)	20.359	20.638	20.238
EMUL/CMD	20.708	20.972	20.632
Processor Util.			
TOTAL (H)	173.53	174.22	171.83
TOTAL	173.00	174.00	172.00
UTIL/PROC (H)	86.76	87.11	85.92
UTIL/PROC	86.50	87.00	86.00
TVR(H)	1.62	1.59	1.62
TVR	1.59	1.57	1.59
Storage			
WKSET (V)	81	67	82
PGBLPGS	40767	40699	40733
PGBLPGS/USER	27.5	27.5	27.5
FREEPGS	3968	3974	3964
FREE UTIL	0.93	0.93	0.93
SHRPGS	1093	1110	1117
Paging			
READS/SEC	466	370	474
WRITES/SEC	325	253	331
PAGE/CMD	15.028	11.771	15.379
XSTOR IN/SEC	0	0	0
XSTOR OUT/SEC	0	0	0
XSTOR/CMD	0.000	0.000	0.000
FAST CLR/CMD	5.661	5.744	5.655

R/O Data	Minidisk	DIRCONTROL w/ D.S.	Minidisk
# Shared FSTs	1 OF 4	4 OF 4	1 OF 4
MINIDISK CACHING	YES	YES	NO
CONTROL UNIT CACHING	YES	YES	NO
RELEASE	ESA 1.1	ESA 1.1	ESA 1.1
RUN ID	L24F1483	L24F1481	L24F1482
Environment			
REAL STORAGE	192M	192M	192M
EXP. STORAGE	64M	64M	64M
USERS	1480	1480	1480
VTAMs	1	1	1
VSCSs	1	1	1
PROCESSORS	2	2	2
I/O			
VIO RATE	369	338	355
VIO/CMD	7.010	6.386	6.782
MDC READS	288	255	246
MDC WRITES	97	97	96
MDC MODS	55	56	55
MDC HIT RATIO	0.87	0.86	0.85
PRIVOPs			
PRIVOP/CMD	30.038	30.559	29.473
DIAG/CMD	22.547	21.848	22.195
DIAG 08/CMD	0.760	0.737	0.726
DIAG 10/CMD	0.000	0.000	0.000
DIAG 58/CMD	1.235	1.228	1.242
DIAG 98/CMD	0.532	0.548	0.516
DIAG A4/CMD	2.641	2.173	2.617
DIAG A8/CMD	1.843	1.700	1.662
DIAG 214/CMD	11.456	11.374	11.310
SIE/CMD	71.301	70.871	70.723
SIE INTCPT/CMD	52.049	52.444	50.920
FREE TOTL/CMD	96.435	94.828	96.399
VTAM Machines			
WKSET (V)	1108	942	1262
TOT CPU/CMD (V)	4.6335	4.7235	4.4894
CP CPU/CMD (V)	2.3220	2.3722	2.2500
VIRT CPU/CMD (V)	2.3115	2.3512	2.2394
DIAG 98/CMD (V)	0.540	0.554	0.527
SFS Servers			
WKSET (V)	1360	1850	1373
TOT CPU/CMD (V)	4.1163	4.1881	4.1286
CP CPU/CMD (V)	1.9632	1.9734	1.9422
VIRT CPU/CMD (V)	2.1531	2.2148	2.1863
FP REQ/CMD(Q)	1.338	1.390	1.333
IO/CMD (Q)	1.842	1.815	1.799
IO TIME/CMD (Q)	0.052	0.046	0.057
SFS TIME/CMD (Q)	0.149	0.129	0.167
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM			

Table 36. SFS Usage of VM Data Spaces - XC Mode (9121-480)

9221-170 / XC Mode Users

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9221-170
- STORAGE:
          RSTOR:   48M
          XSTOR:   16M   All reserved for MDC

- DASD:
            PACK NAME  TYPE
   - SYSTEM: PRF05      3380
             PRF01      3380
             PRFRES     3380
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	3 - 3380-03	2	2	2	4	5
3380-D	1 - 3380-03	1	0	1	1	1

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
                 CONTROLLER  NUMBER  LINESPEED
                 3088-02       1       4.5M
```

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR:  BACTRIAN
- CMS BLOCKSIZE:     4K
- USER VM SIZE:      2M
- USER CMS MODE:     XC
- USER RELSHARE:     100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON
RWSERVE1 (R/W)	SFS	32M/XC	1500	QUICKDSP ON
ROSERVE1 (R/O)	SFS	32M/XC	1500	QUICKDSP ON
CRRRECOV	CRR	32M/XC	100	NONE

4) MEASUREMENT DISCUSSION

All users are run in XC mode. All environments measured have the read/write data in SFS directories but vary in the location of the read-only data. Paging to expanded store is not performed in these environments. Expanded storage is used for minidisk caching. No cache control units exist in the DASD configurations. Both environments are tuned as described in "Recommended 9221 Tuning" on page 223.

H14F0242 R/O Data on minidisks. One of the four minidisks has its FSTs in a shared segment. This case is intended to represent a typical usage of read-only minidisks.

H14M0241 R/O data in SFS DIRCONTROL directories in data spaces. These directories reside in a separate file pool.

The 9221-170 environment, like the 9121-480, does not have expanded storage for paging and therefore benefited more than the 9021-720 from having shared FSTs. When comparing the data spaces environment to minidisks, there was a decrease in paging (PAGE/CMD) of 16% and a 7% improvement in external response time (AVG LAST(T)).

R/O Data // Shared FSTs RELEASE RUN ID	Minidisk 1 of 4 ESA 1.1 H14F0242	DIRCONTROL w/ D.S. 4 of 4 ESA 1.1 H14M0241
Environment		
REAL STORAGE	**48M**	**48M**
EXP. STORAGE	**16M**	**16M**
USERS	**240**	**240**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
Response Time		
TRIV INT	0.176	0.160
NONTRIV INT	1.037	0.986
TOT INT	0.782	0.738
TOT INT ADJ	0.674	0.632
AVG FIRST (T)	0.310	0.290
AVG LAST (T)	0.710	0.660
Throughput		
AVG THINK (T)	27.90	28.02
ETR	7.27	7.20
ETR (T)	8.43	8.41
ETR RATIO	0.862	0.856
ITR (H)	9.94	9.97
ITR	8.57	8.54
EMUL ITR	13.70	13.50
ITRR (H)	1.000	1.003
ITRR	1.000	0.996
Proc. Usage		
PBT/CMD (H)	100.610	100.342
PBT/CMD	100.778	101.064
CP/CMD (H)	43.587	42.843
CP/CMD	37.940	38.048
EMUL/CMD (H)	57.023	57.499
EMUL/CMD	62.838	63.016
Processor Util.		
TOTAL (H)	84.86	84.39
TOTAL	85.00	85.00
UTIL/PROC (H)	84.86	84.39
UTIL/PROC	85.00	85.00
TVR(H)	1.76	1.75
TVR	1.60	1.60
Storage		
WKSET (V)	81	67
PGBLPGS	9638	9611
PGBLPGS/USER	40.2	40.0
FREEPGS	732	744
FREE UTIL	0.87	0.86
SHRPGS	1091	1091
Paging		
READS/SEC	57	47
WRITES/SEC	46	39
PAGE/CMD	12.212	10.225
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.572	5.588

R/O Data # Shared FSTs RELEASE RUN ID	Minidisk 1 of 4 ESA 1.1 H14F0242	DIRCONTROL w/ D.S. 4 of 4 ESA 1.1 H14M0241
Environment		
REAL STORAGE	**48M**	**48M**
EXP. STORAGE	**16M**	**16M**
USERS	**240**	**240**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
I/O		
VIO RATE	57	52
VIO/CMD	6.758	6.183
MDC READS	47	42
MDC WRITES	16	16
MDC MODS	11	11
MDC HIT RATIO	0.89	0.89
PRIVOPs		
PRIVOP/CMD	22.734	22.629
DIAG/CMD	23.060	22.461
DIAG 08/CMD	0.711	0.594
DIAG 10/CMD	0.000	0.000
DIAG 58/CMD	1.186	1.189
DIAG 98/CMD	2.371	2.616
DIAG A4/CMD	2.490	2.140
DIAG A8/CMD	1.778	1.546
DIAG 214/CMD	10.908	10.701
SIE/CMD	67.936	67.772
SIE INTCPT/CMD	51.632	52.185
FREE TOTL/CMD	104.809	103.323
VTAM Machines		
WKSET (V)	190	190
TOT CPU/CMD (V)	18.2532	18.8424
CP CPU/CMD (V)	8.2056	8.4304
VIRT CPU/CMD (V)	10.0477	10.4120
DIAG 98/CMD (V)	2.480	2.716
SFS Servers		
WKSET (V)	397	579
TOT CPU/CMD (V)	12.9280	12.7296
CP CPU/CMD (V)	6.9999	6.8854
VIRT CPU/CMD (V)	5.9281	5.8442
FP REQ/CMD(Q)	1.339	1.375
IO/CMD (Q)	2.144	2.087
IO TIME/CMD (Q)	0.037	0.035
SFS TIME/CMD (Q)	0.050	0.049
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM		

Table 37. SFS Usage of VM Data Spaces - XC Mode (9221-170)

9021-720 - Read/Write Data on Minidisks

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-720
- STORAGE:
   - RSTOR:   512M
   - XSTOR:   2G

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              PSPT01     3380-D
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
		NUMBER OF PACKS				
3380-K	4 - 3990-02	0	0	0	0	16
3380-A	10 - 3880-03	20	8	12	20	0
3380-D	3 - 3880-03	0	0	0	20	0

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	22	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR:  BACTRIAN
- CMS BLOCKSIZE:     4K
- USER VM SIZE:      2M
- USER CMS MODE:     XA/XC
- USER RELSHARE:     100
```

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
SERVE6 (R/O)	SFS	32M/XC	1500	QUICKDSP ON
CRRSERVA	CRR	32M/XC	100	NONE

4) MEASUREMENT DISCUSSION

For measurements shown in this section, the read/write data is on minidisks but the location of the read-only data is varied as follows:

Y64R5865 R/O data on minidisks. One of four minidisks has its FSTs in a shared segment. Minidisk caching is in effect. User virtual machines are run in XA mode.

Y64M5861 R/O data in SFS DIRCONTROL directories in data spaces. User virtual machines are run in XC mode.

Comparing read-only data in SFS directories in data spaces to typical R/O minidisk usage showed similar external response time (AVG LAST (T)). Having all four sets of FSTs in shared memory reduced paging (XSTOR/CMD) by 20%. The slight decrease in internal throughput rate (ITR (H)) is due to the combined effects of going from XA mode to XC mode and SFS usage of data spaces.

R/O DATA USER MODE # SHARED FSTs RELEASE RUN ID	MINIDISK XA 1 OF 4 ESA 1.1 Y64R5865	DIRCONTROL w/ D.S. XC 4 OF 4 ESA 1.1 Y64M5861
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**5860**	**5860**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
Response Time		
TRIV INT	0.056	0.057
NONTRIV INT	0.303	0.314
TOT INT	0.193	0.199
TOT INT ADJ	0.224	0.232
AVG FIRST (T)	0.397	0.410
AVG LAST (T)	0.527	0.550
Throughput		
AVG THINK (T)	25.19	26.22
ETR	238.56	239.53
ETR (T)	205.97	205.07
ETR RATIO	1.158	1.168
ITR (H)	236.35	234.81
ITR	45.71	45.87
EMUL ITR	69.86	69.33
ITRR (H)	1.000	0.994
ITRR	1.000	1.003
Proc. Usage		
PBT/CMD (H)	25.386	25.552
PBT/CMD	25.344	25.504
CP/CMD (H)	9.268	9.066
CP/CMD	8.739	8.631
EMUL/CMD (H)	16.116	16.482
EMUL/CMD	16.604	16.872
Processor Util.		
TOTAL (H)	522.88	524.00
TOTAL	522.00	523.00
UTIL/PROC (H)	87.15	87.33
UTIL/PROC	87.00	87.17
TVR(H)	1.58	1.55
TVR	1.53	1.51
Storage		
WKSET (V)	58	48
PGBLPGS	104K	103K
PGBLPGS/USER	18.2	18.0
FREEPGS	14349	15061
FREE UTIL	0.96	0.93
SHRPGS	1186	1331
Paging		
READS/SEC	332	272
WRITES/SEC	186	126
PAGE/CMD	2.515	1.941
XSTOR IN/SEC	1053	845
XSTOR OUT/SEC	1298	1032
XSTOR/CMD	11.414	9.153
FAST CLR/CMD	5.651	5.710

R/O DATA USER MODE # SHARED FSTs RELEASE RUN ID	MINIDISK XA 1 OF 4 ESA 1.1 Y64R5865	DIRCONTROL w/ D.S. XC 4 OF 4 ESA 1.1 Y64M5861
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**5860**	**5860**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
I/O		
VIO RATE	1751	1625
VIO/CMD	8.501	7.924
MDC READS	1223	1062
MDC WRITES	550	535
MDC MODS	463	456
MDC HIT RATIO	0.93	0.93
PRIVOPs		
PRIVOP/CMD	20.075	20.146
DIAG/CMD	23.603	22.955
DIAG 08/CMD	0.719	0.722
DIAG 10/CMD	0.015	0.015
DIAG 58/CMD	1.219	1.229
DIAG 98/CMD	0.291	0.297
DIAG A4/CMD	4.083	3.540
DIAG A8/CMD	1.893	1.877
DIAG 214/CMD	12.371	12.269
SIE/CMD	54.688	54.928
SIE INTCPT/CMD	35.547	36.252
FREE TOTL/CMD	84.517	84.888
VTAM Machines		
WKSET (V)	1791	1734
TOT CPU/CMD (V)	3.8211	3.8204
CP CPU/CMD (V)	1.8823	1.8788
VIRT CPU/CMD (V)	1.9388	1.9416
DIAG 98/CMD (V)	0.295	0.299
SFS Servers		
WKSET (V)	na	767
TOT CPU/CMD (V)	na	0.0529
CP CPU/CMD (V)	na	0.0388
VIRT CPU/CMD (V)	na	0.0917
FP REQ/CMD(Q)	na	0.016
IO/CMD (Q)	na	0.0
IO TIME/CMD (Q)	na	0.0
SFS TIME/CMD (Q)	na	0.000
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM		

Table 38. SFS Usage of VM Data Spaces - Read/Write Data on Minidisks

3990-3 DASD Fast Write

DASD fast write (DFW) is a 3990 feature which decreases write response time by immediately returning channel end and device end when a write hit occurs. The 3990 controller then processes the write when the DASD is available. The data is saved in nonvolatile storage (NVS), thereby eliminating the possibility of data loss even if a power failure occurs.

A write hit occurs when the DASD record being written is found in the control unit cache. If the DASD record is not found in the cache (write miss), the control unit writes the record out to DASD while simultaneously placing the record into the cache. Once the DASD write has completed, channel end and device end are returned and the I/O request is complete. The control unit then stages all following records on that track into the cache. This means that subsequent writes to these record locations will result in write hits so long as these records remain in the cache.

There are several software settings which have to be on to activate the DASD fast write function: SET CACHE, SET NVS, and SET DASDFW. When the measurement data table indicates that DFW is ON, the run had SET CACHE ON for the 3390 DASD devices/subsystem, SET NVS SUBSYSTEM ON, and SET DASDFW ON.

For additional information on DASD fast write, see "Related Publications" on page 343 for a list of related publications.

9021-580 / Minidisk

This study investigates the effects of 3390-3 DASD fast write on CMS intensive performance.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-580
- STORAGE:
   - RSTOR:   512M
   - XSTOR:   1G

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3390-2	2 - 3990-J03	6	4	8	14/6	0

(NUMBER OF PACKS: PAGE, SPOOL, TDISK, USER, SERVER)

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	2	44	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:   100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON

4) MEASUREMENT DISCUSSION

The DASD were evenly distributed between two 3990-3 control units. Each 3990 contained a 64M cache and a 4M cache for nonvolatile storage (NVS). There were four logical paths to each DASD volume and the DASD I/O activity was distributed across 14 channels.

For all measurements, the user minidisk volumes were made eligible for read caching (the default). The spool volumes were eligible for read caching for all measurements except Y34R309E. In all measurements with DASD fast write on, the user minidisk and spool volumes were additionally made eligible for NVS and DASD fast write.

In all measurements, the page and tdisk volumes were made ineligible for both read caching and write caching (via DASD fast write). The page volumes were made ineligible because the page volumes are marginal candidates for use with either read caching or write caching. This is because paging I/O is already done very efficiently with many pages typically being read or written in a single I/O operation. The tdisk volumes were made ineligible because overall performance was better that way (discussed below).

In all measurements, all user and system CMS minidisks were made eligible for minidisk caching (the default).

Two pairs of measurements were obtained: one with 14 user minidisk volumes and one with 6 user minidisk volumes. This was done in order to evaluate DASD fast write at two different levels of I/O contention.

This discussion will first take a look at what effect DASD fast write had on the performance of the user minidisk volumes and the spool volumes. It will then examine how DASD fast write affected the performance of the system as a whole.

A number of key DASD performance indicators, averaged over the user minidisk volumes, are summarized in the following table:

DFW STATUS MINIDISK VOLUMES RELEASE RUN ID	DFW OFF 14 ESA 1.1 Y34R309E	DFW ON 14 ESA 1.1 Y34R309G	DFW OFF 6 ESA 1.1 Y34R309H	DFW ON 6 ESA 1.1 Y34R309F
Rate (total)	123.2	123.2	118.2	123.0
Rate	8.8	8.8	19.7	20.5
Pct Busy	19.2	10.8	50.2	27.5
Serv	21.9	12.2	25.5	13.8
Resp	28.4	13.8	57.8	20.0
Pct Read	19	19	17	19
Pct Read Hits	58	61	61	62
Pct DFW Hits	na	83	na	83
Pct DeStge	na	12	na	12
Pct DFW Bypass	na	0	na	0

Table 39. DFW Measurements - User Minidisk Volumes

The first group of metrics is from the CACHE_DASD_BY_CONFIG VMPRF report while the second group is taken from the DASD_BY_CONFIG_EF report (new to VMPRF 1.2.1). The meaning of each of these performance indicators is briefly described below:

Rate (total) - I/O rate (per second) summed over all devices in this group.

Rate - I/O rate (per second) to the device.

Pct Busy - Device utilization.

Serv - DASD service time (in milliseconds). This is the sum of pending, connect, and disconnect time.

Resp - DASD response time (msec). This includes DASD service time plus time in queue waiting to start the I/O operation.

Pct Read - The percentage of all I/O operations that are reads. (For these measurements, the remaining I/Os are all writes.)

Pct Read Hits - The percentage of all read I/Os that resulted in a read hit.

Pct DFW Hits - The percentage of all write I/Os that resulted in a write hit.

Pct DeStge - The number of destages divided by the total number of I/Os issued to the device (or control unit) times 100. Destage refers to the movement of updated records from the read cache to the device. With DASD fast write, this can occur asynchronous to the write operation that caused these records to be placed in the cache.

Pct DFW Bypass - The percentage of write requests that were forced (due to NVS constraints) to write directly to DASD.

Because the minidisk volumes were divided evenly between the two 3990 control units and because each minidisk volume had about the same amount of activity, the two sets of minidisk volumes showed very similar average performance characteristics and responded to DASD fast write in an equivalent manner.

As shown in the above table, DASD fast write resulted in substantial decreases in device utilization and DASD response time in both the unconstrained (14 volumes) and constrained (6 volumes) cases. With 14 minidisk volumes, device utilization decreased by 44% and DASD response time decreased by 51%. With

6 minidisk volumes, device utilization decreased by 45% and DASD response time decreased by 65%.

Almost all of the decrease in device service time was due to a decrease in disconnect time. For example, for the 14 minidisk volume case, disconnect time decreased from 17.5 msec to 7.9 msec with the use of DASD fast write. Pending and connect time were essentially unchanged at 0.3 msec and (about) 4.0 msec respectively.

One reason why DASD fast write had such a large impact was that 81% of all I/Os to the minidisk volumes are write requests--all of which are eligible for DASD fast write. This write percentage is so high because many of the read requests are being satisfied out of the minidisk cache in the processor's expanded storage. (The VMPRF data shows that if minidisk caching were turned off, only 60% of the I/Os to the user minidisk volumes would be writes.)

Although there was a substantial decrease in DASD response time, DASD fast write (combined with read caching) will often result in much lower DASD response times than those shown here. For example, the OfficeVision results (see "9021-580 / OfficeVision" on page 188) show average DASD response time dropping to 6.7 milliseconds when DASD fast write is enabled.

What appears to be happening in these FS7B measurements is that the control units are being stressed by the minidisk formatting activity that is occurring on the tdisk volumes. The FS7B workload includes a significant amount of temporary minidisk formatting. In the measured configuration, formatting is going on at the rate of about one 3380 cylinder per second. The OfficeVision IOB workload does have temporary minidisk usage, but the temporary minidisks are formatted during run stabilization and are not reformatted thereafter.

An additional measurement (not shown) was obtained with the tdisk volumes also made eligible for DASD fast write. This did not perform as well because the high level of format activity flooded the NVS (as evidenced by a significant number of DFW bypasses). Although tdisk volumes can be good candidates for DASD fast write, these results suggest that some caution is in order. One approach would be to make the tdisk volumes eligible and monitor Pct DFW Bypass. If it stays near zero, all is well but if it goes above (say) 5%, make the tdisk volumes ineligible and see if DASD performance is better that way.

The following table summarizes how DASD fast write affected the spool volumes:

DFW STATUS READ CACHE STATUS MINIDISK VOLUMES RELEASE RUN ID	DFW OFF OFF 14 ESA 1.1 Y34R309E	DFW ON ON 14 ESA 1.1 Y34R309G	DFW OFF ON 6 ESA 1.1 Y34R309H	DFW ON ON 6 ESA 1.1 Y34R309F
Rate (total)	42.8	47.2	42.0	47.2
Rate	10.7	11.8	10.5	11.8
Pct Busy	16.6	7.7	10.0	7.8
Serv	14.5	4.8	8.7	6.1
Resp	14.5	4.8	8.7	6.1
Pct Read	56	52	51	57
Pct Read Hits	0	95	95	96
Pct DFW Hits	na	91	na	92
Pct DeStge	na	9	na	0
Pct DFW Bypass	na	9	na	0

Table 40. DFW Measurements - User Spool Volumes

Because Y34R309E did not have read caching in effect, the results from the first pair of measurements (14 minidisk volumes) show how the combination of read caching and DASD fast write benefits spool volume performance. Average device utilization decreased by 54% while average DASD response time decreased by 67%. The results from the second pair (6 minidisk volumes) show how DASD fast write alone benefits spool volume performance. Device utilization decreased by 22% while DASD response time decreased by 30%. Read caching and DASD fast write both had large effects on spool volume performance because the spool volume I/Os are split fairly evenly between read and write activity.

The results show quite high DASD fast write hit ratios for the spool volumes. However, DASD fast write decreased DASD response time for the spool volumes to a much lesser extent than it decreased DASD response time for the minidisk volumes. Considering the 6 minidisk case, DASD response time decreased 2.6 msec for the spool volumes and 37.8 msec for the minidisk volumes. This is because 1) the spool volumes are much less I/O constrained than the minidisk volumes and 2) even without DASD fast write, the spool volumes had a very low DASD response time (8.7 msec). This is because reads are very fast due to the very high read hit percentage (95%) and because CP does a excellent job of managing the spool volume allocations and I/O processing so as to minimize seek and latency time.

As with the minidisk volumes, nearly all of the decrease in spool volume service time was due to a decrease in disconnect time.

These minidisk volume and spool volume performance benefits translated into improved total system responsiveness. As shown in Table 41 on page 186, external response time (AVG LAST (T)) decreased by about 6% in both the unconstrained (14 minidisk volumes) and constrained (6 minidisk volumes) cases. Internal response times (TOT INT ADJ) decreased by 11% in the unconstrained case and 14% in the constrained case. Most of the improvements resulted from the use of DASD fast write for the minidisk volumes. This is because there were 2.6 times as many minidisk I/Os as spool I/Os and because the amount of DASD response time reduction was much greater for the minidisk volumes.

The degree to which DASD fast write improves system response time depends upon a number of factors. Some of the more influential ones are:

1. I/O intensiveness of the workload

   FS7B is not highly I/O intensive. In the measured environment, it does approximately 2.3 non-paging DASD I/Os per command (about 2.8 additional I/Os are handled via minidisk caching). Other workloads that do more I/Os per command have the potential to show larger total system impacts. For example, the OfficeVision IOB workload does about 8.4 DASD I/Os per command (see "9021-580 / OfficeVision" on page 188 for results).

   For the four FS7B measurements shown, page I/Os per command (PAGE/CMD) ranged from 2.6 to 3.4. Therefore, in the measured environment, over half of all DASD I/Os were to the page volumes, which were ineligible for DASD fast write.

2. the percentage of total system DASD I/Os to which DASD fast write applies

   For FS7B in the measured configuration, 45% of all DASD I/Os are writes to DFW-eligible devices and about 38% of all DASD I/Os are writes that experience DASD fast write hits.

3. the amount of DASD response time reduction per DFW hit

   As discussed above, although these FS7B measurements showed large decreases, workloads with less temp disk formatting activity may show even larger reductions.

These results illustrate that, with DASD fast write, the number of DASD actuators can be reduced while preserving or even improving system responsiveness. Going from 14 minidisk volumes without DASD fast write (Y34R309E) to 6 minidisk volumes with DASD fast write (Y34R309F) resulted in a net 3.6% decrease in external response time.

This evaluation was done with the FS7B0R workload which does not have any SFS usage (all files are on minidisks). Our expectation is that DASD fast write would benefit the equivalent SFS usage workloads (FS7B35R and FS7BMAXR) by about the same extent. In addition, the use of DASD fast write for the SFS log minidisks and catalog minidisks may be of value in reducing the probability of rollbacks due to deadlock.

In this study, the three system volumes were left on 3380s behind a 3880 control unit and were therefore not included in this DASD fast write evaluation. However, volumes such as these that contain active CP data areas can be good candidates for use with DASD fast write.

These results demonstrate that DASD fast write can be quite effective at reducing I/O subsystem contention and improving DASD responsiveness. The best improvements were observed for minidisk volumes because spool volume I/O is already optimized by CP. Finally, the results suggest that caution should be exercised when using DASD fast write with tdisk volumes.

DFW STATUS MINIDISK VOLUMES RELEASE RUN ID	DFW OFF 14 ESA 1.1 Y34R309E	DFW ON 14 ESA 1.1 Y34R309G	DFW OFF 6 ESA 1.1 Y34R309H	DFW ON 6 ESA 1.1 Y34R309F
Environment				
REAL STORAGE	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**1024M**	**1024M**	**1024M**	**1024M**
USERS	**3090**	**3090**	**3090**	**3090**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**
PROCESSORS	**3**	**3**	**3**	**3**
Response Time				
TRIV INT	0.041	0.038	0.046	0.041
NONTRIV INT	0.279	0.248	0.315	0.270
TOT INT	0.186	0.166	0.212	0.181
TOT INT ADJ	0.200	0.179	0.227	0.195
AVG FIRST (T)	0.250	0.243	0.240	0.243
AVG LAST (T)	0.387	0.363	0.397	0.373
Throughput				
AVG THINK (T)	25.80	25.61	25.56	25.78
ETR	118.07	119.07	117.93	118.64
ETR (T)	109.92	110.69	110.31	110.37
ETR RATIO	1.074	1.076	1.069	1.075
ITR (H)	124.89	125.10	124.61	124.38
ITR	44.75	44.88	44.47	44.59
EMUL ITR	66.43	66.69	66.29	66.18
ITRR (H)	1.000	1.002	0.998	0.996
ITRR	1.000	1.003	0.994	0.996
Proc. Usage				
PBT/CMD (H)	24.021	23.980	24.075	24.120
PBT/CMD	24.017	23.940	24.024	24.101
CP/CMD (H)	8.195	8.187	8.280	8.219
CP/CMD	7.824	7.769	7.887	7.883
EMUL/CMD (H)	15.821	15.789	15.791	15.896
EMUL/CMD	16.193	16.171	16.137	16.219
Processor Util.				
TOTAL (H)	264.04	265.45	265.56	266.20
TOTAL	264.00	265.00	265.00	266.00
UTIL/PROC (H)	88.01	88.48	88.52	88.73
UTIL/PROC	88.00	88.33	88.33	88.67
TVR(H)	1.52	1.52	1.52	1.52
TVR	1.48	1.48	1.49	1.49
Storage				
WKSET (V)	62	62	62	62
PGBLPGS	50053	49368	49223	50061
PGBLPGS/USER	16.2	16.0	15.9	16.2
FREEPGS	8300	8875	8904	8309
FREE UTIL	0.96	0.98	0.98	0.96
SHRPGS	982	960	1056	996
Paging				
READS/SEC	204	231	248	247
WRITES/SEC	82	80	130	100
PAGE/CMD	2.602	2.810	3.427	3.144
XSTOR IN/SEC	505	510	479	495
XSTOR OUT/SEC	618	620	643	620
XSTOR/CMD	10.216	10.208	10.172	10.103
FAST CLR/CMD	5.486	5.411	5.485	5.473

DFW STATUS MINIDISK VOLUMES RELEASE RUN ID	DFW OFF 14 ESA 1.1 Y34R309E	DFW ON 14 ESA 1.1 Y34R309G	DFW OFF 6 ESA 1.1 Y34R309H	DFW ON 6 ESA 1.1 Y34R309F
Environment				
REAL STORAGE	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**1024M**	**1024M**	**1024M**	**1024M**
USERS	**3090**	**3090**	**3090**	**3090**
VTAMs	**1**	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**	**2**
PROCESSORS	**3**	**3**	**3**	**3**
I/O				
VIO RATE	911	916	916	923
VIO/CMD	8.288	8.275	8.304	8.363
MDC READS	604	610	614	610
MDC WRITES	292	295	283	291
MDC MODS	240	243	238	239
MDC HIT RATIO	0.92	0.92	0.93	0.92
PRIVOPs				
PRIVOP/CMD	21.348	21.301	21.424	21.201
DIAG/CMD	11.240	11.207	11.243	11.308
DIAG 08/CMD	0.746	0.741	0.752	0.752
DIAG 10/CMD	0.009	0.009	0.009	0.009
DIAG 58/CMD	1.255	1.247	1.251	1.250
DIAG 98/CMD	0.418	0.416	0.399	0.408
DIAG A4/CMD	3.839	3.839	3.826	3.851
DIAG A8/CMD	1.910	1.897	1.931	1.957
DIAG 214/CMD	na	na	na	na
SIE/CMD	58.806	58.621	58.628	58.650
SIE INTCPT/CMD	39.988	39.863	39.867	39.882
FREE TOTL/CMD	83.978	85.389	84.003	83.992
VTAM Machines				
WKSET (V)	954	996	956	962
TOT CPU/CMD (V)	3.8462	3.8445	3.8429	3.8155
CP CPU/CMD (V)	1.9155	1.9172	1.9088	1.8977
VIRT CPU/CMD (V)	1.9307	1.9272	1.9340	1.9178
DIAG 98/CMD (V)	0.427	0.420	0.407	0.412
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM				

Table 41. DFW Measurements - CMS Intensive Workload

9021-580 / OfficeVision

The following 9021-580 measurements document the performance advantage for the OfficeVision environment of using the 3990 DASD fast write (DFW) support introduced in this release.

1) WORKLOAD: IOB V2.1

2) HARDWARE CONFIGURATION

```
- PROCESSOR:
     9021-580:    9021-720  CPU 0, 1, 2 only
- STORAGE:
   - RSTOR:   256 M
   - XSTOR:   1024 M

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              PSPT01     3380-D
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-D	3 - 3880-3	10	0	0	0	0
3390-2	2 - 3990-L03	0	6	4	14	8

(NUMBER OF PACKS: PAGE, SPOOL, TDISK, USER, SERVER)

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	36	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: IOB
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:    100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
PRODBM	OV/VM	16M/XA	10000	QUICKDSP ON
PROCAL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMAIL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMBX00 - 50	OV/VM	16M/XA	10000	QUICKDSP ON; IBCENTRL=Y

4) MEASUREMENT DISCUSSION

The IBM Office Benchmark (IOB) makes use of service machines that perform synchronous I/O. Improving the DASD service times for these servers allows them the potential of handling higher transaction rates, as well as improving their service time. Thirty-two 3390-2 DASD devices were connected to two 3990-3 control units. Each control unit contained a 128M cache with 4M of nonvolatile storage (NVS). The DASD data shown below was extracted from VMPRF 1.2.1 reports CACHE_DASD_BY_CONFIG, DASD_BY_CONFIG_EF and UCLASS_STATES. The meaning of each of these performance indicators is as follows:

Rate (total) - I/O rate (per second) summed over all devices in this group.

Pct Busy - Device utilization.

Serv - DASD service time (in milliseconds). This is the sum of pending, connect, and disconnect time.

Resp - DASD response time (msec). This includes DASD service time plus time in queue waiting to start the I/O operation.

Pct Read - The percentage of all I/O operations that are reads. (For these measurements, the remaining I/Os are all writes.)

Pct Read Hits - The percentage of all read I/Os that resulted in a read hit.

Pct DFW Hits - The percentage of all write I/Os that resulted in a write hit.

Pct DeStge - The number of destages divided by the total number of I/Os issued to the device (or control unit) times 100. Destage refers to the movement of updated records from the read cache to the device. With DASD fast write, this can occur asynchronous to the write operation that caused these records to be placed in the cache.

Pct DFW Bypass - The percentage of write requests that were forced (due to NVS constraints) to write directly to DASD.

CPU Secs - The total processor time expended, in seconds.

Pct True Dormnt - The percentage of elapsed time a user is in true dormant state.

Pct True Non-Dormnt - The percentage of elapsed time a user is in true non-dormant state.

Pct Test Idle - The percentage of true non-dormant time that a user was found in test idle.

To get an estimate as to the improvement in capacity due to reduced DASD response time, the following values are calculated:

CPU Msec/Sec - The average CPU usage (msec) consumed each second by the server machine(s) being analyzed (1000 * CPU Secs / Elapsed time).

DASD Resp Msec/Sec - The average DASD response time (msec) consumed each second by the server machine(s) being analyzed (Rate (total) * Resp).

Busy Msec/Sec - The sum of the times the server machine(s) are using CPU or performing DASD I/O (CPU Msec/Sec + DASD Resp Msec/Sec).

Active Msec/Sec - The sum of the time the server machine(s) are active. This variable is calculated by summing the inactive time and subtracting the results from 1000 milliseconds. The inactive time consists of true dormant time (1000 * Pct True Dormnt / 100) plus the test idle time (1000 * (Pct True Non-Dormnt / 100) * (Pct Test Idle / 100)).

The following table shows the data for the one DASD volume used exclusively for the Calendar service machine's calendar files. With DASD fast write set on, the DASD response time improved by 43%. One reason that the DASD response time did not improve as much as other DASD described in this section is the low Pct Read hits. This is due to the random selection of user calendars that are being reviewed and updated.

The Calendar machine capacity improvement, due to its synchronous nature, is bounded by the time it is consuming resources (Busy Msec/Sec) and the time the machine is considered to be active (Active Msec/Sec). The difference between these numbers represents time spent waiting for the availability of resources and depends largely on other events occurring outside this virtual machine. Therefore, for this workload the improvement in resource consumption ranged between 29% and 39% for the Calendar machine. Putting this in terms of capacity, the Calendar machine should be able to support from 1.4 to 1.6 times as much activity through the use of DASD fast write.

DFW STATUS **RELEASE** **RUN ID**	**DFW OFF** **ESA 1.1** **Y34V3251**	**DFW ON** **ESA 1.1** **Y34V3252**	**DELTA**	**PERCENT**
Rate (total)	8.4	8.4	0.0	0
Pct Busy	23.6	13.2	-10.4	-44
Serv	27.9	15.8	-12.1	-43
Resp	27.9	15.8	-12.1	-43
Pct Read	26	26	-	-
Pct Read Hits	44	35	-	-
Pct DFW Hits	na	95	-	-
Pct DeStge	na	12	-	-
Pct DFW Bypass	na	0	-	-
CPU Secs	48	49	-	-
Pct True Dormnt	26.2	30.2	-	-
Pct True Non-Dormnt	73.8	69.8	-	-
Pct Test Idle	58.0	68.6	-	-
CPU Msec/Sec	26.7	27.2	0.5	2
DASD Resp Msec/Sec	234.4	132.7	-101.7	-43
Busy Msec/Sec	261.1	159.9	-101.2	-39
Active Msec/Sec	310	219	-91	-29

Table 42. DFW Measurements - Calendar Minidisk Volume

The following table shows the data for the one DASD volume used exclusively for the documents stored by the Database machine. With DASD fast write set on, the DASD response time improved by 64%. Using the same rationale as described above for the Calendar machine, the improvement in resource consumption for this workload ranged between 45% and 54% for the Database machine. Putting this in terms of capacity, the Database machine should be able to support from 1.8 to 2.2 times as much activity through the use of DASD fast write.

DFW STATUS RELEASE RUN ID	DFW OFF ESA 1.1 Y34V3251	DFW ON ESA 1.1 Y34V3252	DELTA	PERCENT
Rate (total)	13.5	13.1	-0.4	-3
Pct Busy	26.0	9.2	-16.8	-65
Serv	19.3	7.0	-12.3	-64
Resp	19.3	7.0	-12.3	-64
Pct Read	15	15	-	-
Pct Read Hits	70	53	-	-
Pct DFW Hits	na	100	-	-
Pct DeStge	na	4	-	-
Pct DFW Bypass	na	0	-	-
CPU Secs	89	93	-	-
Pct True Dormnt	17.0	25.2	-	-
Pct True Non-Dormnt	83.0	74.8	-	-
Pct Test Idle	52.1	70.7	-	-
CPU Msec/Sec	49.4	51.7	2.3	5
DASD Resp Msec/Sec	260.6	91.7	-169.0	-65
Busy Msec/Sec	309.0	143.4	-165.6	-54
Active Msec/Sec	398	219	-179	-45

Table 43. DFW Measurements - Database Minidisk Volume

The following table shows the data for the six DASD volumes used exclusively for the Mailbox service machines, containing the user in-basket disks and the server A-disks. With DASD fast write set on, the DASD response time improved by 77% for this workload. With the use of multiple servers, the I/O done and CPU used by one server can be asynchronous to that done by another server. For this reason the calculation used for active time would not be applicable in this case. However, it would be expected to be somewhat less than the 71% improvement shown in Busy Msec/Sec. Another way to look at capacity is the number of Mailbox machines required to support a given workload. This can be done by dividing Busy Msec/Sec by the maximum number of Msec/Sec that can be consumed by one server (1000). Without DASD fast write, this configuration required at least three Mailbox server machines (2507.7 / 1000). With DASD fast write on, this requirement is reduced to one (720.8 / 1000). It should be noted that these are theoretical minimums. It is not practical to run a server at or near 100% utilization. There are many other factors that need to be considered when determining the number of Mailbox machines, but this data does imply that fewer Mailbox machines will be required with DASD fast write active.

DFW STATUS RELEASE RUN ID	DFW OFF ESA 1.1 Y34V3251	DFW ON ESA 1.1 Y34V3252	DELTA	PERCENT
Rate (total)	67.4	66.2	-1.2	-2
Pct Busy	29.6	8.3	-21.3	-72
Serv	25.7	7.3	-18.4	-72
Resp	34.1	7.7	-26.4	-77
Pct Read	8	8	-	-
Pct Read Hits	63	57	-	-
Pct DFW Hits	na	98	-	-
Pct DeStge	na	13	-	-
Pct DFW Bypass	na	0	-	-
CPU Secs	377	380	-	-
Pct True Dormnt	78.5	89.8	-	-
Pct True Non-Dormnt	21.5	10.2	-	-
Pct Test Idle	28.8	66.9	-	-
CPU Msec/Sec	209.4	211.1	1.7	1
DASD Resp Msec/Sec	2298.3	509.7	-1788.6	-78
Busy Msec/Sec	2507.7	720.8	-1786.9	-71

Table 44. DFW Measurements - Mailbox Minidisk Volumes

The following table shows the data for the thirty-two DASD fast write volumes used for the entire measurement, with only paging and some system volumes on non DASD fast write units. With DASD fast write set on, the DASD response time improved by 75% for this workload.

DFW STATUS RELEASE RUN ID	DFW OFF ESA 1.1 Y34V3251	DFW ON ESA 1.1 Y34V3252	DELTA	PERCENT
Rate (total)	481.0	478.3	-2.7	-1
Pct Busy	25.7	8.3	-17.4	-68
Serv	19.1	6.1	-13.0	-68
Resp	26.3	6.7	-19.6	-75
Pct Read	23	23	-	-
Pct Read Hits	83	85	-	-
Pct DFW Hits	na	95	-	-
Pct DeStge	na	9	-	-
Pct DFW Bypass	na	0	-	-
DASD Resp Msec/Sec	12650.3	3204.6	-9445.7	-75

Table 45. DFW Measurements - All DASD Fast Write Volumes

The following table shows the overall system effects of the measurements discussed above. The internal response time (TOT INT ADJ) was reduced by 0.153 seconds (25%) with no change in CPU usage. The expected improvement to internal response time, due to DASD fast write, was calculated by dividing the DASD response time improvement (DASD Resp Msec/Sec from above) by the external commands per second (ETR (T) below). This calculation indicates that the response time should be improved by 0.156 seconds, which closely matches the observed improvement.

The third column (Y34V3351) shows the effect of increasing the number of users with DASD fast write on. This measurement still showed good performance, but does not illustrate the types of capacity improvement discussed above, because the CPU was the constrained resource. DASD fast write only improved the I/O subsystem. Had the system been constrained on I/O, the entire system capacity would have been improved by a greater extent.

DFW STATUS RELEASE RUN ID	DFW OFF ESA 1.1 Y34V3251	DFW ON ESA 1.1 Y34V3252	DFW ON ESA 1.1 Y34V3351
Environment			
REAL STORAGE	**256M**	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**	**1024M**
USERS	**3244**	**3248**	**3339**
VTAMs	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**
PROCESSORS	**3**	**3**	**3**
Response Time			
TRIV INT	0.036	0.033	0.046
NONTRIV INT	0.669	0.500	0.682
TOT INT	0.506	0.380	0.512
TOT INT ADJ	0.614	0.461	0.639
AVG FIRST (T)	0.653	0.527	0.693
AVG LAST (T)	0.953	0.733	0.983
Throughput			
AVG THINK (T)	42.81	43.11	42.48
ETR	73.44	73.50	77.72
ETR (T)	60.52	60.56	62.32
ETR RATIO	1.214	1.214	1.247
ITR (H)	69.55	69.57	69.21
ITR	28.21	28.16	28.76
EMUL ITR	50.66	50.55	51.76
ITRR (H)	1.000	1.000	0.995
ITRR	1.000	0.998	1.019
Proc. Usage			
PBT/CMD (H)	43.136	43.124	43.349
PBT/CMD	43.128	43.098	43.328
CP/CMD (H)	20.446	20.417	20.641
CP/CMD	19.168	18.989	19.257
EMUL/CMD (H)	22.683	22.700	22.701
EMUL/CMD	23.960	24.108	24.071
Processor Util.			
TOTAL (H)	261.05	261.16	270.13
TOTAL	261.00	261.00	270.00
UTIL/PROC (H)	87.02	87.05	90.04
UTIL/PROC	87.00	87.00	90.00
TVR(H)	1.90	1.90	1.91
TVR	1.80	1.79	1.80
Storage			
WKSET (V)	53	54	53
PGBLPGS	52189	52223	51927
PGBLPGS/USER	16.1	16.1	15.6
FREEPGS	6949	6950	7163
FREE UTIL	0.96	0.96	0.96
SHRPGS	1150	1158	1165
Paging			
READS/SEC	13	15	20
WRITES/SEC	30	32	39
PAGE/CMD	0.711	0.776	0.947
XSTOR IN/SEC	1003	966	1082
XSTOR OUT/SEC	1096	1061	1190
XSTOR/CMD	34.684	33.471	36.460
FAST CLR/CMD	15.615	15.505	15.502

DFW STATUS RELEASE RUN ID	DFW OFF ESA 1.1 Y34V3251	DFW ON ESA 1.1 Y34V3252	DFW ON ESA 1.1 Y34V3351
Environment			
REAL STORAGE	**256M**	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**	**1024M**
USERS	**3244**	**3248**	**3339**
VTAMs	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**
PROCESSORS	**3**	**3**	**3**
I/O			
VIO RATE	1353	1360	1401
VIO/CMD	22.357	22.457	22.482
MDC READS	721	723	745
MDC WRITES	536	534	550
MDC MODS	454	452	466
MDC HIT RATIO	0.89	0.89	0.89
PRIVOPs			
PRIVOP/CMD	24.511	24.750	24.594
DIAG/CMD	89.752	90.683	90.292
DIAG 08/CMD	10.427	10.684	10.688
DIAG 10/CMD	0.000	0.000	0.000
DIAG 58/CMD	1.851	1.849	1.861
DIAG 98/CMD	0.694	0.710	0.674
DIAG A4/CMD	11.831	11.840	11.859
DIAG A8/CMD	6.444	6.489	6.499
DIAG 214/CMD	35.659	35.931	35.705
SIE/CMD	139.430	140.390	141.425
SIE INTCPT/CMD	92.024	94.061	93.341
FREE TOTL/CMD	236.889	236.725	246.487
VTAM Machines			
WKSET (V)	1568	1602	1607
TOT CPU/CMD (V)	4.8287	4.8960	4.8320
CP CPU/CMD (V)	2.4235	2.4755	2.4249
VIRT CPU/CMD (V)	2.4052	2.4205	2.4071
DIAG 98/CMD (V)	0.705	0.711	0.683
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM			

Table 46. The Effects of 3990-3 DASD Fast Write in an OfficeVision Environment

CMS Pipelines

Two sets of instruction traces and a multi-user benchmark run were obtained. The first set of traces was a group of 6 commands traced on both PRPQ 1.1.6 and VM/ESA 1.1 CMS Pipelines. The second set was a group of “equivalent function” traces that compared REXX/EXECIO/XEDIT statements to their equivalent pipelines specification. This set was used to determine how well CMS Pipelines performs relative to existing methods for given functions.

In addition to the trace data collected, system performance comparisons were made between the REXX/EXECIO/XEDIT and CMS Pipelines functions. A multi-user benchmark test was made on a 9021-580 that compared 3000 users running the REXX/EXECIO/XEDIT versions of the “equivalent functions” to 3000 users running the CMS Pipelines versions.

The traces and the multi-user benchmark were run on VM/ESA 1.1. All I/O was to minidisk or spool for all commands traced or issued in the multi-user benchmark. CMS Pipelines was not installed in a shared segment for the PRPQ vs. VM traces; however, it was installed in a shared segment for the REXX/EXECIO/XEDIT vs. CMS Pipelines tests. See “CMS Pipelines” on page 290 for more information on the virtual machine configuration of the users.

Comparison to PRPQ 1.1.6 CMS Pipelines

Trace data was collected for each command using the CP TRACE command. See “Measurement Methodology” on page 243 for a description of the methodology used to collect and reduce the trace data.

The following 6 commands were traced for PRPQ 1.1.6 and VM/ESA 1.1 CMS Pipelines:

1. ***PIPE CMS Q DISK | > QUERY DISK A***

   This command is an example of issuing CMS commands from CMS Pipelines and stream I/O to disk.

2. ***PIPE < NATHAN NAMES A| CONSOLE***

   This command is an example of stream I/O to disk and console.

3. ***PIPE CP Q N|SPLIT ,|STRIP |LOCATE /- DSC/|COUNT LINES|SPEC *-* 1 /Users disconnected/ NEXT| CONSOLE***

   This command is an example of CP commands from CMS Pipelines and a number of various CMS Pipelines filters.

4. ***PIPE (end \) < NATHAN NAMES A|c:LOCATE /:nic/|SPEC 24-* 1|JOIN 2 / /| LITERAL Ids:|CONSOLE \c:|SPEC 24-* 1|JOIN 2 / /|LITERAL NAMES:| > NA OUT A***

   This command is an example of multi-stream CMS Pipelines and more filters.

5. ***PIPE < NATHAN NAMES A|SPEC 1 A|CONSOLE***

   This command issues a CMS Pipelines specific error message

6. ***PIPE LITERAL A RECORD | DUP 9 | FANIN | COUNT LINES | CONSOLE***

   This command has more filters.

The following table gives estimated virtual CPU time, pathlength and privileged operations use. The CPU times shown here are actually time estimates calculated by the STARS reduction tool (see "Measurement Methodology" on page 243) to account for the differences in instruction mix between the two traces. These CPU times do not represent actual system CPU time. These times are provided mainly to account for the differences in opcode usage between PRPQ 1.1.6 and VM/ESA 1.1 CMS Pipelines and should only be used in comparison between the CMS Pipelines traces.

	Estimated Virtual CPU Time (msecs)				Virtual Pathlength				Privops			
No.	PRPQ	VM	Delta	Pct.	PRPQ	VM	Delta	Pct.	PRPQ	VM	Delta	Pct.
1	43.823	43.865	0.042	0.10%	53507	53425	-82	-0.15%	254	254	0	0%
2	38.322	38.335	0.023	0.06%	53544	53487	-57	-0.11%	386	386	0	0%
3	31.211	31.149	-0.062	-0.20%	46176	45895	-281	-0.61%	104	104	0	0%
4	68.807	68.770	-0.037	-0.05%	96884	96530	-354	-0.37%	231	231	0	0%
5	25.752	25.215	-0.537	-2.09%	34953	34844	-104	-0.30%	168	156	-12	-7.14%
6	21.342	21.330	-0.012	-0.06%	30252	30136	-116	-0.38%	103	103	0	0%
AVG	38.210	38.111	-0.099	-0.26%	52553	52386	-167	-0.32%	208	206	2	-0.96%

Table 47. PRPQ 1.1.6 vs. VM/ESA 1.1 CMS Pipelines Pathlength and Privop Counts

The following table breaks out the storage use by each trace into shared and non-shared pages. It further breaks out non-shared pages into references to the CMS Pipelines code and other data areas in non-shared storage. This is done since references to the CMS Pipelines code would be shared page references if CMS Pipelines was put into a shared segment.

	PRPQ 1.1.6				ESA 1.1				DELTA			
		Non-Shared				Non-Shared				Non-Shared		
No.	Total	Pipes	Other	Shared	Total	Pipes	Other	Shared	Total	Pipes	Other	Shared
1	126	22	40	64	126	22	40	64	0	0	0	0
2	93	20	32	41	93	20	32	41	0	0	0	0
3	93	28	33	32	90	25	33	32	-3	-3	0	0
4	115	27	44	44	116	28	44	44	1	1	0	0
5	89	24	30	35	95	23	30	42	6	-1	0	7
6	82	21	29	32	83	22	29	32	1	1	0	0
AVG	100	24	35	41	101	23	35	43	1 (1%)	0	0	1 (3%)

Table 48. PRPQ 1.1.6 vs. VM/ESA 1.1 CMS Pipelines Storage Use

Pathlength/CPU Time: For the particular commands traced, VM/ESA 1.1 CMS Pipelines exhibited a 0.26% decrease in estimated CPU time and a 0.32% decrease in pathlength. The reader may note that in the first two traces the pathlength was shorter for VM/ESA 1.1 CMS Pipelines, yet the estimated CPU time was greater. This is due to more expensive instructions being issued by VM/ESA 1.1 CMS Pipelines than PRPQ 1.1.6 and illustrates the need to include the CPU times with the pathlength.

Storage Use: VM/ESA 1.1 CMS Pipelines referenced on average one more page per command than PRPQ 1.1.6; however, the second table shows that VM/ESA 1.1 and PRPQ 1.1.6 CMS Pipelines were equivalent in their references to "other" non-shared storage. This "other" non-shared storage includes CMS low non-shared storage, CMS control blocks unique to that user (such as the page allocation table), and any CMS Pipelines related data areas. The CMS Pipelines non-shared storage refers to that area of the user's virtual machine where the CMS Pipelines code is loaded. Since CMS Pipelines can be put into a shared segment, all non-shared page references would then be restricted to the “other” category and illustrates that PRPQ 1.1.6 and VM/ESA 1.1 CMS Pipelines are equivalent in user working storage requirements.

Privileged Operations: VM/ESA 1.1 CMS Pipelines was equivalent to PRPQ 1.1.6 CMS Pipelines with the exception of trace 5 where VM/ESA 1.1 pipes issued fewer privileged operations.

REXX/EXECIO/XEDIT versus CMS Pipelines

Twelve “functions” were coded in REXX/EXECIO/XEDIT and with CMS Pipelines. The functions chosen are ones that can be written with one or two pipeline specifications. This was done since it is felt that this would be the most common (and perhaps the most beneficial) use of CMS Pipelines. That is, application developers would use CMS Pipelines to replace what would have required several (possibly many) lines of REXX/EXECIO code (most likely involving loops) or the use of XEDIT macros. Most I/O performed by these functions is to disk or console since this was felt to be the most common case; however, two tests used spool I/O.

What was not tested were applications that could only be coded in CMS Pipelines. It was felt that in these cases the application developer would have to use CMS Pipelines. Additionally, applications that could much more easily be coded using the previously available methods over CMS Pipelines were not tested. In these cases the application developer would probably choose to continue with the previous methods. Thus, this study focused on situations where CMS Pipelines saves coding time and effort over the previously available methods.

These twelve “functions” were constructed to test as many different features of CMS Pipelines and coding possibilities with REXX/EXECIO/XEDIT as possible. They are by no means all inclusive nor do they test every possible application of CMS Pipelines, REXX, EXECIO or XEDIT. For example, no punch I/O, print I/O, or fullscreen applications are included in any of the tests. They were constructed only to provide some idea of how CMS Pipelines compares to previously available coding methods. Each of the following features of CMS Pipelines (and their REXX/EXECIO/XEDIT equivalents) are included in at least one test: issuing CP commands, issuing CMS commands, file I/O, reader I/O, a subset of filters to manipulate data, multi-stream pipelines and filters in REXX.

The twelve “functions” are described below and the corresponding REXX and CMS Pipelines equivalents are provided in “CMS Pipelines” on page 290.

1. Read in an entire small (30 lines in this case) file and store in a REXX STEM variable.

2. Issue a CP command and store the response in a REXX STEM.

3. Issue the CMS Q DISK command and store the response in a REXX STEM.

4. Write a single line of output to a file (includes opening and closing of file).

5. Issue the CMS Q DISK command searching for the mode of a particular disk.

6. Read in a small file (30 lines) and output to the screen.

7. Issue the CP Q NAMES command and count the number of disconnected users.

8. Read in a small file (30 lines) and output some selected lines to the screen while writing others to a disk file.

   Two REXX/EXECIO versions were tested in this case. The first version used a “one line at a time” file read and output approach while the second version read in the entire file before processing, performed some work and output the results. These test cases are referred to as 8a and 8b respectively.

9. Read in only a certain portion of a small file (30 line file) and store in a REXX STEM.

   Two REXX/EXECIO/XEDIT versions were tested in this case. The first version used XEDIT macros to extract the selected portion of the file while the second version version used EXECIO to read in the selected portion of the file. These test cases are referred to as 9a and 9b respectively.

10. Read in certain lines of a small file (meeting some search criteria) and output to the console.

    Two different CMS Pipelines versions were tested in this case. Both versions used a filter coded in REXX but used different filters in accomplishing the same task. These test cases are referred to as 10a and 10b respectively.

11. Read in a reader file punched with the NOH option and store in a REXX stem variable.

12. Read in a reader file punched with the NOH option, reorganizing the columns of the records and writing the output to disk.

Traces

The following three tables summarize the results of the REXX vs. CMS Pipelines traces. The first table shows estimated CPU time and pathlength. The second and third tables show storage use (pages referenced) and privileged operations respectively. The third table also lists the unassisted individual privops with counts in parentheses that had a non-zero delta between the REXX and CMS Pipelines traces. The values here are calculated by the same method used in "Comparison to PRPQ 1.1.6 CMS Pipelines" on page 196, AVGB shown for CMS Pipelines and REXX is the average of the best case test versions (some tests cases have two versions) while AVGT is the average for all versions of all test cases.

	Estimated CPU Time				Pathlength			
No.	REXX	Pipelines	Delta	Pct.	REXX	Pipelines	Delta	Pct.
1	54.050	77.594	23.544	43.6%	66669	97878	31209	46.8%
2	21.995	28.779	6.784	30.8%	25698	35325	9627	37.5%
3	55.868	50.355	-5.513	-9.9%	71014	62412	-8602	-12.1%
4	26.960	36.311	9.351	34.7%	29787	44247	14460	48.5%
5	73.220	56.908	-16.312	-22.3%	94381	73658	-20723	-22.0%
6	136.745	67.830	-68.915	-50.4%	170180	94569	-75611	-44.4%
7	44.084	39.076	-5.000	-11.3%	55887	53294	-2593	-4.6%
8a 8b	569.688 376.928	102.707	-466.981 -274.221	-82.0% -72.7%	765756 487359	144829	-620927 -342530	-81.1% -70.3%
9a 9b	109.796 64.090	43.242	-66.554 -20.848	-60.6% -32.5%	133664 82850	57500	-76164 -25350	-57.0% -30.6%
10a 10b	594.933	124.929 218.153	-470.004 -376.780	-79.0% -63.3%	790902	167672 293284	-623230 -497618	-78.8% -62.9%
11	53.963	74.791	20.828	38.6%	64777	97729	32952	50.9%
12	230.169	76.370	-153.799	-66.8%	290388	99837	-190551	-65.6%
AVGB AVGT	144.417 172.321	64.907 76.696	-79.509 -95.625	-55.1% -55.5%	185824 223522	85745 101710	-100079 -121812	-53.9% -54.5%

Table 49. REXX/EXECIO/XEDIT vs. CMS Pipelines Pathlength and CPU Time

	Total				Shared				Non-Shared			
No.	REXX	Pipes	Delta	Pct.	REXX	Pipes	Delta	Pct.	REXX	Pipes	Delta	Pct.
1	98	115	17	17.3%	61	81	20	32.8%	37	34	-3	-8.1%
2	99	104	5	5.1%	62	73	9	14.5%	37	31	-6	-16.2%
3	111	126	15	13.5%	79	96	17	21.5%	32	30	-2	-6.3%
4	106	119	13	12.3%	64	78	14	21.9%	42	41	-1	-2.4%
5	127	135	8	6.3%	89	104	15	16.9%	38	31	-7	-18.4%
6	101	110	9	8.9%	63	78	15	23.8%	38	32	-6	-15.8%
7	104	115	11	10.6%	66	83	17	25.8%	38	32	-6	-15.8%
8a 8b	117 115	134	17 19	14.5% 16.5%	71 68	89	18 21	25.4% 30.9%	46 47	45	-1 -2	-2.2% -4.3%
9a 9b	185 112	109	-76 -3	-41.1% -2.7%	129 74	78	-46 4	-35.7% 5.4%	56 38	31	-25 -7	-44.6% -18.4%
10a 10b	222	157 156	-65 -66	-29.3% -29.7%	147	108 108	-39 -39	-26.5% -26.5%	75	49 49	-26 -26	-34.7% -34.7%
11	99	111	12	12.1%	62	79	17	27.4%	37	32	-5	-13.5%
12	111	128	17	15.3%	67	84	17	25.4%	44	44	0	0%
AVGB AVGT	117 122	122 125	5 3	4.3% 2.5%	75 79	86 88	11 9	14.7% 11.4%	42 43	36 37	-6 -6	-14.3% -14.0%

Table 50. REXX/EXECIO/XEDIT vs. CMS Pipelines Storage Use

No.	Total Privops				Unassisted Privops			
	REXX	Pipelines	Delta	Pct.	REXX	Pipelines	Delta	Pct.
1	331	467	136	41.1%	22	31	9	40.1%
2	151	199	48	31.8%	20	22	2	10.0%
3	264	339	75	28.4%	19	21	2	10.5%
4	168	233	65	38.7%	29	29	0	0.0%
5	276	305	29	10.5%	19	23	4	21.1%
6	901	730	-171	-19.0%	142	104	-38	-26.8%
7	195	192	-3	-1.5%	23	24	1	4.4%
8a 8b	1726 608	364	-1362 -244	-78.9% -40.1%	53 55	52	-1 -4	-1.9% -7.3%
9a 9b	443 380	229	-214 -151	-48.3% -39.7%	55 22	24	-31 2	-56.4% 9.1%
10a 10b	1826	760 1342	-1066 -484	-58.4% -26.5%	122	69 67	-53 -55	-43.4% -45.1%
11	358	375	17	4.7%	23	26	3	13.0%
12	629	236	-393	-62.5%	35	37	2	5.7%
AVGB AVGT	507 590	369 444	-138 -146	-27.2% -24.7%	44 46	39 41	-5 -5	-11.4% -10.9%

Table 51. REXX/EXECIO/XEDIT vs. CMS Pipelines Privileged Operations Use

Pathlength: As shown in the table, the average CPU time decreased 55.1% and the average pathlength decreased 53.9% when implementing these twelve equivalent functions with CMS Pipelines rather than REXX/EXECIO/XEDIT. The estimated CPU times are provided here to account for the instruction mix of the traces along with the pathlength. The CPU time will be correspondingly greater if a given trace has more privileged operations and/or expensive instructions. These results indicated that test cases 1, 2, 4 and 11 favored the REXX/EXECIO/XEDIT implementations over CMS Pipelines while the others all favored CMS Pipelines.

Privileged Operations: The results show that CMS Pipelines used an average of 27.2% fewer privileged operations per command than the REXX/EXECIO/XEDIT test cases. The privops used in these test cases were all assisted privops on ESA hardware with the exception of the DIAG, SSCH and TSCH instructions. The counts for the SSCH and TSCH instructions were the same for both CMS Pipelines and REXX/EXECIO/XEDIT in all the test cases. CMS Pipelines issued an average of 11.4% fewer unassisted privops than REXX/EXECIO/XEDIT per test case, all being DIAG instructions.

Storage: On average, CMS Pipelines referenced five more pages per trace than the REXX equivalents; however, these extra page references were to shared storage and CMS Pipelines actually required six less non-shared pages on average. This is significant since there is only one copy of shared code in real storage while there must be a separate copy of non-shared storage for each user. CMS Pipelines used the same or less non-shared storage in every test case, resulting in a smaller demand for real storage. This would not be true if CMS Pipelines was not installed in a shared segment as it was for these test cases. In the PRPQ 1.1.6 vs. VM/ESA 1.1 CMS Pipelines comparison, there were over 20 pages referenced in the CMS Pipelines code for each trace. Assuming that the number of pages referenced would be similar for these test cases, one can see that CMS Pipelines would use an average of over 14 more pages of non-shared storage.

Examination of the Test Cases: The test cases favoring REXX/EXECIO were reading of a file into storage, issuing a CP command and saving the results in storage, writing a line to a file, and reading a reader spool file into storage. The common feature of all these test cases is that the function can be handled via one EXECIO statement versus one pipe specification. Thus, the simple act of reading in a file or issuing CP commands was faster with EXECIO than CMS Pipelines. It was also just as easy to code the function with EXECIO as it was with CMS Pipelines. It appears that a user simply replacing EXECIO statements with Pipe statements in an EXEC will suffer performance-wise and not gain much in productivity.

The test cases favoring CMS Pipelines were issuing a CMS command storing results, determining mode of a disk in a search order, reading in a file and outputting to screen, determining the number of disconnected users, read in a file and output certain lines to screen with others output to disk, read in only a selected portion of a file, read in a file and output to screen lines matching a search criteria, and read in a reader spool file reorganizing the columns of data. The common feature of all these test cases was that the implementation of the function in REXX/EXECIO/XEDIT required a few or many lines of code while the pipelines implementation still only required one pipe specification (with exception of test case 10). Test case 10 was an example of a CMS Pipelines filter written in REXX and this case was still faster than using the REXX/XEDIT implementation. In these cases, the overhead of interpreting and running REXX loops, EXECIO or XEDIT was greater than the overhead of building the pipeline and running the data through the various CMS Pipelines filters. These results not only show the power of CMS Pipelines to replace many REXX statements (often involving loops) with a pipeline specification (with filters) but also show that the performance of the EXEC will improve, sometimes by as much as 2-5 times faster depending on the function and implementation.

Multi-user Benchmark

This benchmark consisted of running 3000 identical remote users on a 9021-580. These users had virtual machine configurations identical to the ones used for the traces. Two tests were completed: the first test had all 3000 users repeatedly executing the REXX/EXECIO/XEDIT versions of the twelve test cases presented in the above section while the second test consisted of all 3000 users repeatedly executing the CMS Pipelines versions of the twelve test cases.

System Configuration: The following is a description of the environment used to test VM/ESA 1.1 CMS Pipelines vs. REXX/EXECIO/XEDIT multi-user benchmark.

1) WORKLOAD: CMS PIPELINES

2) HARDWARE CONFIGURATION

```
- PROCESSOR:   9021-580
- STORAGE:
   - RSTOR:        256M
   - XSTOR:          1G

- DASD:
              PACK NAME   TYPE
   - SYSTEM:  RESPAK     3380-A
              SRVPAK     3380-A
              ESAP01     3380-A
              ESAOV1     3380-A
              ESAOV2     3380-A
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3390-A	2 - 3990-3	4	5	5	10	0
3380-A	1 - 3880-2	0	0	0	5	0

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CTCA	NUMBER	CHANNEL SPEED
3088	1	4.5M

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:    64

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON

4) MEASUREMENT DISCUSSION

The following table provides a view of the system performance with 3000 users running the REXX equivalents vs. 3000 users running the CMS Pipelines equivalents. Each user executed all twelve test cases (multiple versions where applicable) repeatedly throughout the whole run. The first run shows the results where all test cases were the REXX/EXECIO/XEDIT versions while the second run shows the results for the CMS Pipelines versions.

The multi-user benchmark test was run to further reinforce the results of the traces and was used to show the difference in total system performance between using REXX/EXECIO/XEDIT vs. CMS Pipelines. As shown here, the external response time (AVG LAST (T)) decreased by 0.07 seconds (14.7%) and internal throughput rate (ITR (H)) increased by 16.2%. The traces show a large decrease in virtual pathlength which is shown here as a 29.0% decrease in emulation time per command (EMUL/CMD (H)). Note that there were a few other commands that the users had to run with the test cases which explains why the decrease in EMUL/CMD was not as great as in the traces. The traces had shown an average of 6 fewer non-shared pages per test case for the CMS Pipelines equivalents; the multi-user run exhibited an average decrease in working set size (WKSET (V)) of 8 pages. The system paging per command (PAGE/CMD + XSTOR/CMD) correspondingly showed a decrease of 27.0%. The MDC hit ratios were near 100% for both runs due to the small number of files used in the run relative to the amount of XSTOR on the system and due to the uniformity of the workload. The virtual I/Os per command (VIO/CMD) was approximately the same in both runs, as would be expected since both the REXX/EXECIO/XEDIT and CMS Pipelines equivalents of the test cases issued the same amount of I/O. The traces also showed a decrease in DIAG instructions per test case and the same number of SSCH and TSCH instructions per trace for the REXX/EXECIO/XEDIT and CMS Pipelines equivalents. This is reflected in the multi-user benchmark as a decrease in DIAG/CMD and an approximately equal number of unassisted privops per command (PRIVOP/CMD).

IMPLEMENTATION RELEASE	REXX/EXECIO/XEDIT ESA 1.1	CMS Pipelines ESA 1.1
Environment		
REAL STORAGE	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**
USERS	**3000**	**3000**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**3**	**3**
Response Time		
TRIV INT	0.169	0.141
NONTRIV INT	1.039	1.022
TOT INT	0.305	0.277
TOT INT ADJ	0.241	0.210
AVG FIRST (T)	0.310	0.260
AVG LAST (T)	0.430	0.367
Throughput		
AVG THINK (T)	25.34	25.36
ETR	97.24	96.46
ETR (T)	123.08	127.29
ETR RATIO	0.790	0.758
ITR (H)	149.64	173.84
ITR	39.43	43.90
EMUL ITR	78.66	105.60
ITRR (H)	1.000	1.162
ITRR	1.000	1.113
Proc. Usage		
PBT/CMD (H)	20.049	17.258
PBT/CMD	20.068	17.283
CP/CMD (H)	10.343	10.362
CP/CMD	9.993	10.134
EMUL/CMD (H)	9.703	6.892
EMUL/CMD	10.075	7.149
Processor Util.		
TOTAL (H)	246.76	219.67
TOTAL	247.00	220.00
UTIL/PROC (H)	82.25	73.22
UTIL/PROC	82.33	73.33
TVR(H)	2.07	2.50
TVR	1.99	2.42
Storage		
WKSET (V)	47	39
PGBLPGS	45000	45072
PGBLPGS/USER	15.0	15.0
FREEPGS	8612	8675
FREE UTIL	0.95	0.94
SHRPGS	623	592
Paging		
READS/SEC	0	0
WRITES/SEC	8	9
PAGE/CMD	0.065	0.071
XSTOR IN/SEC	590	444
XSTOR OUT/SEC	590	444
XSTOR/CMD	9.587	6.976
FAST CLR/CMD	4.655	4.030

IMPLEMENTATION RELEASE	REXX/EXECIO/XEDIT ESA 1.1	CMS Pipelines ESA 1.1
Environment		
REAL STORAGE	**256M**	**256M**
EXP. STORAGE	**1024M**	**1024M**
USERS	**3000**	**3000**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**3**	**3**
I/O		
VIO RATE	724	747
VIO/CMD	5.882	5.868
MDC READS	122	116
MDC WRITES	191	159
MDC MODS	191	159
MDC HIT RATIO	1.00	0.99
PRIVOPs		
PRIVOP/CMD	14.760	14.799
DIAG/CMD	19.168	15.030
DIAG 08/CMD	0.244	0.338
DIAG 10/CMD	0.000	0.000
DIAG 58/CMD	0.000	0.000
DIAG 98/CMD	0.357	0.369
DIAG A4/CMD	1.414	1.257
DIAG A8/CMD	0.463	0.267
DIAG 214/CMD	11.196	9.710
SIE/CMD	43.800	38.785
SIE INTCPT/CMD	29.784	27.925
FREE TOTL/CMD	232.950	249.383
VTAM Machines		
WKSET (V)	1388	1364
TOT CPU/CMD (V)	4.4661	4.5260
CP CPU/CMD (V)	1.0060	0.9951
VIRT CPU/CMD (V)	3.4601	3.5309
DIAG 98/CMD (V)	0.360	0.369

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

Table 52. REXX/EXECIO/XEDIT vs. CMS Pipelines Multi-User Benchmark Test

GCS IPOLL Option

9021-720 / 35% SFS

This section shows the performance effects of specifying IPOLL ON for the FS7B35R CMS intensive workload.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

- PROCESSOR: 9021-720
- STORAGE:
  - RSTOR: 512M
  - XSTOR: 2048M

- DASD:

	PACK NAME	TYPE
- SYSTEM:	PSYS02	3380-A
	PSPT01	3380-D
	WKLD01	3380-D
	WKLD02	3380-D

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	10 - 3880-3	20	8	12	0	0
3380-K	4 - 3990-2	0	0	0	0	16

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	22	56Kb

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2M
- USER CMS MODE: XA
- USER RELSHARE: 100

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
SERVE2	SFS	32M/XA	1500	QUICKDSP ON
SERVE4	SFS	32M/XA	1500	QUICKDSP ON
SERVE7	SFS	32M/XA	1500	QUICKDSP ON
SERVE8	SFS	32M/XA	1500	QUICKDSP ON
CRRSERVA	CRR	16M/XA	100	

4) MEASUREMENT DISCUSSION

The results show that the use of IPOLL ON caused unassisted privops (PRIVOP/CMD) to decrease by 15%. The RTM data (not shown) confirms that all of this decrease results from a decrease in IUCV requests, consistent with the expected effect of IPOLL ON. This decrease in IUCV requests resulted in a 0.55% decrease in CP processing time (CP/CMD (H)) and an overall 0.27% decrease in overall processing requirements (PBT/CMD (H)). The amount of CPU usage improvement was sufficiently small that IPOLL ON had no discernible impact on response time.

IPOLL SETTING RELEASE RUN ID	IPOLL OFF ESA 1.1 Y64F480K	IPOLL ON ESA 1.1 Y64F480V
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**4800**	**4800**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
Response Time		
TRIV INT	0.050	0.048
NONTRIV INT	0.341	0.340
TOT INT	0.221	0.220
TOT INT ADJ	0.245	0.245
AVG FIRST (T)	0.330	0.333
AVG LAST (T)	0.497	0.497
Throughput		
AVG THINK (T)	25.65	25.77
ETR	187.61	187.94
ETR (T)	169.24	168.82
ETR RATIO	1.109	1.113
ITR (H)	202.11	202.65
ITR	37.33	37.56
EMUL ITR	58.91	59.12
ITRR (H)	1.000	1.003
ITRR	1.000	1.006
Proc. Usage		
PBT/CMD (H)	29.687	29.607
PBT/CMD	29.722	29.618
CP/CMD (H)	11.392	11.329
CP/CMD	10.872	10.781
EMUL/CMD (H)	18.289	18.272
EMUL/CMD	18.849	18.837
Processor Util.		
TOTAL (H)	502.41	499.83
TOTAL	503.00	500.00
UTIL/PROC (H)	83.74	83.31
UTIL/PROC	83.83	83.33
TVR(H)	1.62	1.62
TVR	1.58	1.57
Storage		
WKSET (V)	64	64
PGBLPGS	108K	108K
PGBLPGS/USER	23.0	23.0
FREEPGS	12218	12237
FREE UTIL	0.96	0.96
SHRPGS	1362	1349
Paging		
READS/SEC	295	297
WRITES/SEC	159	164
PAGE/CMD	2.683	2.731
XSTOR IN/SEC	1079	1105
XSTOR OUT/SEC	1292	1320
XSTOR/CMD	14.010	14.365
FAST CLR/CMD	5.613	5.615

IPOLL SETTING RELEASE RUN ID	IPOLL OFF ESA 1.1 Y64F480K	IPOLL ON ESA 1.1 Y64F480V
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**4800**	**4800**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
I/O		
VIO RATE	1160	1162
VIO/CMD	6.854	6.883
MDC READS	1037	1031
MDC WRITES	308	307
MDC MODS	240	239
MDC HIT RATIO	0.93	0.93
PRIVOPs		
PRIVOP/CMD	28.852	24.518
DIAG/CMD	21.136	21.133
DIAG 08/CMD	0.727	0.735
DIAG 10/CMD	0.012	0.012
DIAG 58/CMD	1.235	1.244
DIAG 98/CMD	0.313	0.314
DIAG A4/CMD	2.712	2.707
DIAG A8/CMD	1.613	1.653
DIAG 214/CMD	11.540	11.533
SIE/CMD	72.608	66.722
SIE INTCPT/CMD	50.826	45.371
FREE TOTL/CMD	96.811	103.116
VTAM Machines		
WKSET (V)	1496	1492
TOT CPU/CMD (V)	3.9863	3.9194
CP CPU/CMD (V)	1.9882	1.8955
VIRT CPU/CMD (V)	1.9981	2.0239
DIAG 98/CMD (V)	0.318	0.315
SFS Servers		
WKSET (V)	988	953
TOT CPU/CMD (V)	3.7336	3.7285
CP CPU/CMD (V)	1.8012	1.7968
VIRT CPU/CMD (V)	1.9324	1.9317
FP REQ/CMD(Q)	1.347	1.336
IO/CMD (Q)	1.978	1.972
IO TIME/CMD (Q)	0.037	0.037
SFS TIME/CMD (Q)	0.086	0.088
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM		

Table 53. GCS IPOLL Option - CMS Intensive Environment

9021-720 / OfficeVision

The following 9021-720 runs are provided to show the effects of using the new GCS IPOLL option in an OfficeVision environment.

1) WORKLOAD: IOB V2.1

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-720
- STORAGE:
   - RSTOR:   512 M
   - XSTOR:   2048 M

- DASD:
              PACK NAME   TYPE
   - SYSTEM:  PSYS02      3380-A
              PSPT01      3380-D
              WKLD01      3380-D
              WKLD02      3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-D	20 - 3880-3	20	16	12	40	0
3380-D	1 - 3880-G23	0	0	0	0	4
3380-A	3 - 3880-G23	0	0	0	0	12
3380-K	2 - 3990-2	0	0	0	0	16

(NUMBER OF PACKS: PAGE, SPOOL, TDISK, USER, SERVER)

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	36	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: IOB
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:    100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON
PRODBM	OV/VM	16M/XA	10000	QUICKDSP ON
PROCAL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMAIL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMBX00 - 50	OV/VM	16M/XA	10000	QUICKDSP ON; IBCENTRL=Y

4) MEASUREMENT DISCUSSION

Two OfficeVision measurements were compared with the only difference being the GCS IPOLL setting. IPOLL was set on for the VTAM service machine and the two VSCS server machines. An external response time (AVG LAST (T)) improvement of 0.094 seconds (10%) was observed at a cost of 1.0% in CPU time per command (PBT/CMD (H)). The VTAM service machines netted an 8.2% increase in CPU time usage (TOT CPU/CMD (V)), with CP decreasing by 5.5% (CP CPU/CMD (V)) and virtual increasing by 21% (VIRT CPU/CMD (V)).

IPOLL SETTING RELEASE RUN ID	IPOLL OFF ESA 1.1 Y64V620F	IPOLL ON ESA 1.1 Y64V620B
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**6201**	**6200**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
Response Time		
TRIV INT	0.036	0.036
NONTRIV INT	0.403	0.401
TOT INT	0.303	0.302
TOT INT ADJ	0.375	0.371
AVG FIRST (T)	0.627	0.627
AVG LAST (T)	0.927	0.833
Throughput		
AVG THINK (T)	42.50	42.37
ETR	142.20	141.91
ETR (T)	115.02	115.38
ETR RATIO	1.236	1.230
ITR (H)	130.96	129.69
ITR	27.05	26.62
EMUL ITR	50.20	48.79
ITRR (H)	1.000	0.990
ITRR	1.000	0.984
Proc. Usage		
PBT/CMD (H)	45.817	46.264
PBT/CMD	45.731	46.282
CP/CMD (H)	22.595	22.514
CP/CMD	21.040	21.061
EMUL/CMD (H)	23.217	23.746
EMUL/CMD	24.691	25.221
Processor Util.		
TOTAL (H)	526.99	533.80
TOTAL	526.00	534.00
UTIL/PROC (H)	87.83	88.97
UTIL/PROC	87.67	89.00
TVR(H)	1.97	1.95
TVR	1.85	1.84
Storage		
WKSET (V)	53	53
PGBLPGS	108K	108K
PGBLPGS/USER	17.8	17.8
FREEPGS	13129	13014
FREE UTIL	0.97	0.96
SHRPGS	1349	1411
Paging		
READS/SEC	34	30
WRITES/SEC	77	74
PAGE/CMD	0.965	0.901
XSTOR IN/SEC	1640	1659
XSTOR OUT/SEC	1796	1813
XSTOR/CMD	29.873	30.092
FAST CLR/CMD	14.980	14.977

IPOLL SETTING RELEASE RUN ID	IPOLL OFF ESA 1.1 Y64V620F	IPOLL ON ESA 1.1 Y64V620B
Environment		
REAL STORAGE	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**
USERS	**6201**	**6200**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**6**	**6**
I/O		
VIO RATE	2498	2514
VIO/CMD	21.718	21.789
MDC READS	1369	1379
MDC WRITES	979	987
MDC MODS	833	841
MDC HIT RATIO	0.90	0.90
PRIVOPs		
PRIVOP/CMD	22.684	18.145
DIAG/CMD	85.547	85.526
DIAG 08/CMD	9.372	9.430
DIAG 10/CMD	0.000	0.000
DIAG 58/CMD	1.861	1.863
DIAG 98/CMD	0.409	0.416
DIAG A4/CMD	11.624	11.666
DIAG A8/CMD	6.260	6.275
DIAG 214/CMD	35.185	34.989
SIE/CMD	133.542	124.250
SIE INTCPT/CMD	88.138	79.520
FREE TOTL/CMD	222.570	230.750
VTAM Machines		
WKSET (V)	1869	1788
TOT CPU/CMD (V)	4.9288	5.3371
CP CPU/CMD (V)	2.3847	2.2541
VIRT CPU/CMD (V)	2.5440	3.0830
DIAG 98/CMD (V)	0.414	0.418
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 54. The Effects of the GCS IPOLL Setting in an OfficeVision Environment

Inter-System Facility for Communication (ISFC)

The Inter-System Facility for Communications, ISFC, is a Control Program (CP) function which supports transparent, high-performance communications between cooperative applications on VM/ESA Release 1.1 systems and LAN-based programmable workstations running VM Programmable Workstation Communication Services (VM PWSCS).

Figure 9 shows LAN-based workstations communicating with the VM/ESA system through a VM PWSCS domain controller. The domain controller may be connected to the VM/ESA system by an IBM 3088 Multisystem Channel Communication Unit (MCCU). This domain controller workstation acts as a program-to-program communications gateway between the VM/ESA system and the LAN-based workstations, all of which are running VM PWSCS.

The workstations must also be running one of the following operating environments:

- OS/2 Extended Edition
- IBM Personal Computer Disk Operating System (DOS)
- Microsoft Windows
- IBM Advanced Interactive Executive (AIX) for PS/2 (AIX 1.2) or RISC System 6000 (AIX 3.1).

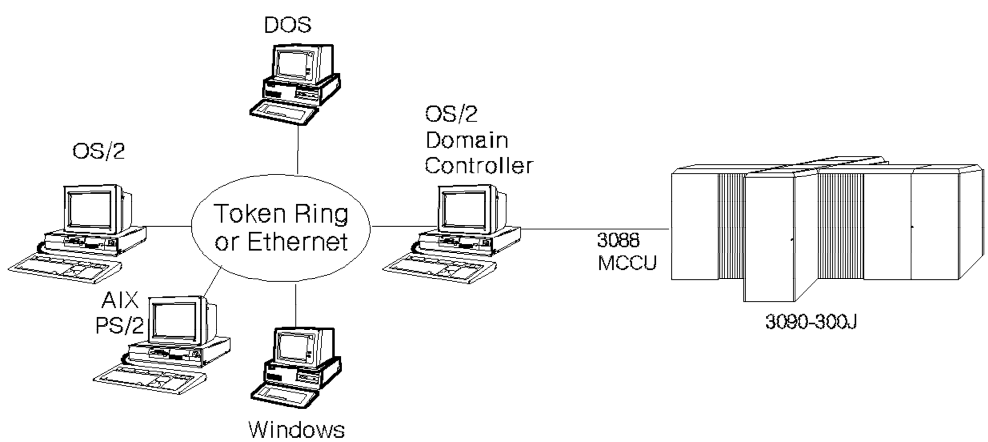

Figure 9. Communications with LAN-based Workstations Running VM PWSCS

Figure 10 shows program-to-program communication within a Communication Services (CS) collection. A user program on a LAN-attached VM PWSCS workstation has a CPI-C program connected to a VM resource manager. The VM resource manager resides in a virtual machine in a VM/ESA system. The VM resource manager in the CMS virtual machine controls access to one or more VM resources.

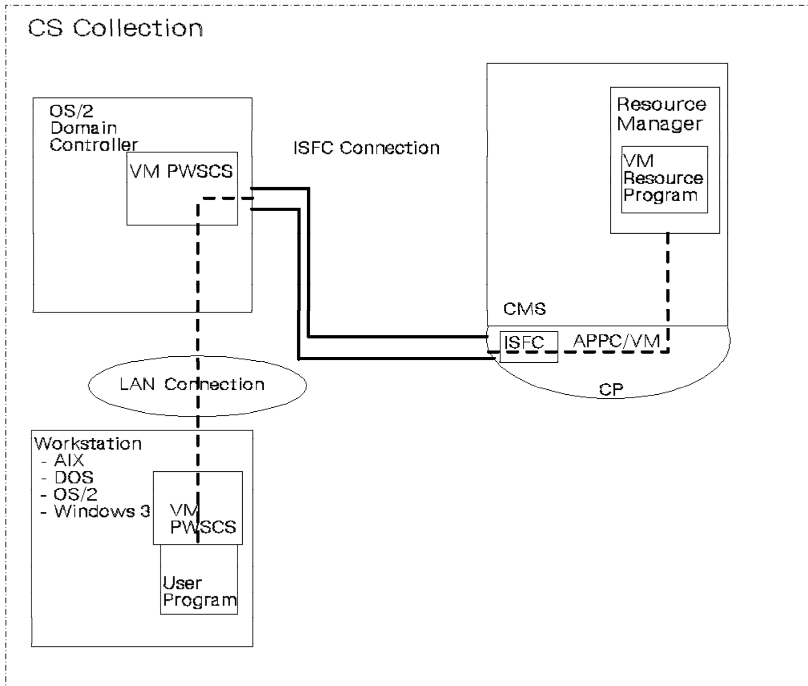

Figure 10. Connectivity within a Communication Services (CS) Collection

VM PWSCS Domain Controller to 3090-300J

The following measurement shows throughput from a VM PWSCS domain controller to a 3090-300J system. The CS collection consists of an OS/2 domain controller and a VM/ESA 1.1 system. The domain controller is equipped with an IBM 370 Channel Adapter /A. This adapter allows the PS/2 to be cabled directly into an IBM 3088 MCCU. These measurements were performed with little background load on the system. The only systems connected into the 3088 were the communications partners themselves.

1) WORKLOAD: INSTVER

2) HARDWARE/SOFTWARE CONFIGURATION

```
- CLIENT PROCESSOR:        PS/2 MODEL 80-121
- OPERATING SYSTEM:        OS/2 1.3 EXTENDED EDITION
- SOFTWARE:                VM PWSCS 1.1 DOMAIN CONTROLLER
- SYSTEM MEMORY:           14M
- FIXED DISK:              115M
- CHANNEL ADAPTER:         370 CHANNEL ADAPTER /A CONFIGURED ABOVE 1M
- PWSCS CONFIGURATION:     DEFAULTS FOR PACING AND BUFFER COUNT


- SERVER PROCESSOR:        3090-300J
- OPERATING SYSTEM:        VM/ESA 1.1
- SOFTWARE:                INTER-SYSTEM FACILITY FOR COMMUNICATIONS
- STORAGE:
   - RSTOR:                256M
   - XSTOR:                1G
- COMMUNICATIONS:
                           CTCA      NUMBER  CHANNEL SPEED
                           3088        1         4.5M
```

3) MEASUREMENT DISCUSSION

The following measurement showed that data throughput increased when large message sizes were utilized. The Common Programming Interface for Communications (CPI-C) architecture utilizes a 16-bit Logical Length (LL) field as part of each send-receive frame that is transmitted over the network. As a consequence, the largest send or receive allowed is about 32000 bytes, since room must also be set aside to accommodate the header information. In this measurement as the size of the message increased toward the 32K limit, the data throughput rate also increased.

There were two main factors which accounted for the throughput increase as the message size increased. The main reason was that the large message sizes made the most efficient use of the Application Programming Interface. Each send incurred a fixed amount of overhead to cross the layers which defined the communications architecture. A large send incurred less overhead since more data was sent per API crossing. There was less header information sent when large messages were used. The ratio of data to header information was greater since fewer header records needed to be sent.

This can be illustrated with a simple example which compares sending 1000 bytes using a message size of 100 bytes to one using a message size of 1000 bytes. Sending ten 100 byte messages involves sending ten header records and ten API crossings. Sending one 1000 byte message involves sending one header record and one API crossing. This explanation is somewhat simplistic in that there are obviously architectural and hardware limitations to the amount of data that can be processed into a single communications packet or frame. When a 32000 byte send is processed, most of the frames that are sent will be complete frames, and only the last frame sent is incomplete. When large messages are sent at the API level, the number of complete frames tends to increase, while the amount of overhead tends to decrease. When complete frames are sent, data throughput tends to increase.

The amount of improvement that can be achieved by increasing message size alone is limited. This is due to the underlying frame packaging going on at the physical hardware level as well as other system constraints. For example, large amounts of resource on the PS/2 are consumed in copying data from memory to the 370 Channel Adapter /A. The limits to system throughput are closely related to CPU and memory speed, rather than just the size of the message used.

MESSAGE SIZE (Bytes)	TIME PER ITERATION (ms)	TRANSFER RATE (Kb/sec)
100	2.50	40.73
1000	3.13	320.51
4000	5.58	717.49
8000	10.99	727.93
12000	16.07	746.97
16000	20.89	765.92
20000	25.85	773.84
24000	30.82	778.84
28000	35.43	790.40
32000	39.91	801.90

Note: These measurements were obtained by running the INSTVER benchmark between an OS/2 domain controller and a 3090-300J system.

Table 55. VM PWSCS Domain Controller to 3090-300J

OS/2 1.3 EE Gateway Attach Measurements

The following measurement was obtained by attaching a user workstation to the collection managed by the OS/2 domain controller. The INSTVER benchmark is run with the message flowing from the user workstation over token ring to the domain controller. The data is packaged at the domain controller, going from token ring to the 370 Channel Adapter /A to the VM/ESA host system.

1) WORKLOAD: INSTVER

2) HARDWARE/SOFTWARE CONFIGURATION

```
- CLIENT PROCESSOR:          PS/2 MODEL 80-121
- OPERATING SYSTEM:          OS/2 1.3 EXTENDED EDITION
- SOFTWARE:                  VM PWSCS 1.1 USER WORKSTATION
- SYSTEM MEMORY:             10M
- FIXED DISK:                115M
- COMMUNICATIONS ADAPTER:    IBM TOKEN RING 16-4 /A RUNNING AT 4 Mbps
- TRANSMIT BUFFER SIZE:      4K
- PWSCS CONFIGURATION:       DEFAULTS FOR PACING AND BUFFER COUNT


- GATEWAY PROCESSOR:         PS/2 MODEL 80-121
- OPERATING SYSTEM:          OS/2 1.3 EXTENDED EDITION
- SOFTWARE:                  VM PWSCS 1.1 DOMAIN CONTROLLER
- SYSTEM MEMORY:             14M
- FIXED DISK:                115M
- CHANNEL ADAPTER:           370 CHANNEL ADAPTER /A CONFIGURED ABOVE 1M
- COMMUNICATIONS ADAPTER:    IBM TOKEN RING 16-4 /A RUNNING AT 4 Mbps
- TRANSMIT BUFFER SIZE:      4K
- PWSCS CONFIGURATION:       DEFAULTS FOR PACING AND BUFFER COUNT


- SERVER PROCESSOR:          3090-300J
- OPERATING SYSTEM:          VM/ESA 1.1
- SOFTWARE:                  INTER-SYSTEM FUNCTION FOR COMMUNICATIONS
- STORAGE:
   - RSTOR:                  256M
   - XSTOR:                  1G
- COMMUNICATIONS:
                             CTCA     NUMBER   CHANNEL SPEED
                             3088       1         4.5M
```

3) MEASUREMENT DISCUSSION

The results of this gateway measurement showed that increasing the message size had a salutary effect on the throughput. The impact of changing the message size was not as pronounced for messages above 4000 bytes. Small messages of 100 bytes incurred greater overhead. The limiting factor was the user workstation's ability to copy data from the application to the token ring, and the domain controller's ability to copy data from token ring to the 370 Channel Adapter /A. The messages tended to spread out along the wire. The token ring, channel adapter, 3088, and VM/ESA system all contained data in various states of processing. When this spacing occurred, the packaging of the messages was no longer optimal. Although a large message was sent from the user workstation, by the time it reached the domain controller and was repackaged, many of the large scale economies were eliminated.

In the following measurement, the token ring was running at 4 Mbps. The transmit buffer size was set to 4 Kb. as recommended earlier.

MESSAGE SIZE (Bytes)	TIME PER ITERATION (ms)	TRANSFER RATE (Kb/sec)
100	4.7	21.4
4000	14.4	278.6
8000	28.3	283.1
12000	41.4	289.9
16000	54.5	293.5
20000	67.9	294.0
24000	83.8	286.5
28000	95.1	294.4
32000	106.6	300.1

Note: Measurements obtained running INSTVER from a user workstation through a domain controller to a VM/ESA 1.1 system.

Table 56. VM PWSCS Gateway Attached Workstation to VM ESA 1.1 System

The following measurement employed a configuration similar to the previous one. The only change that was made was that the token ring data rate was increased to 16 Mbps and the transmit buffer size was increased to 8 Kb. Although the throughput rate plateaus at 360 Kbps, this tends to verify that the token ring on the user workstation was the principle source of the bottleneck. When the token ring adapter rate and buffer size were increased, throughput improved somewhat.

The benefits of sending large messages were not that pronounced. This was because large amounts of system CPU were consumed by the PS/2 in copying from system memory to adapter memory. There was a lot of message packaging going on in the LAN collection. The domain controller must copy data from the token ring to the channel adapter while the user workstation strains to pump data out onto the LAN. The performance bottleneck was the CPU power of the workstations. These I/O operations took a fixed amount of time regardless of the size of the buffer.

MESSAGE SIZE (Bytes)	TIME PER ITERATION (ms)	TRANSFER RATE (Kb/sec)
100	4.16	24.0
1000	5.84	171.2
4000	11.31	353.7
8000	23.78	336.4
12000	34.75	345.3
16000	45.26	353.5
20000	55.54	360.1
24000	65.84	364.5
28000	77.81	359.9
32000	87.91	364.0

Note: These measurements were obtained by running the INSTVER benchmark shipped with VM PWSCS 1.1 for a variety of iteration counts. The PS/2 is running VM PWSCS and is configured as a user workstation. The server is a VM/ESA 1.1 system with ISFC.

Table 57. Gateway Measurements of VM PWSCS 1.1 on a 16 MB/sec LAN

ECKD-Formatted DASD versus CKD-Formatted DASD

These measurements compare the paging performance of ECKD-formatted DASD to the paging performance of CKD-formatted DASD.

3090-300J

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  3090-300J
- STORAGE:
   - RSTOR:   256M
   - XSTOR:   100M

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              PSPT01     3380-D
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-K	1 - 3990-02	4	2	0	0	0
3380-A	6 - 3880-03	0	4	6	8	0
3380-D	2 - 3880-03	0	0	0	12	0

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	44	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:           TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE:    4K
- USER VM SIZE:     2M
- USER CMS MODE:    XA
- USER RELSHARE:    100
```

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON

4) MEASUREMENT DISCUSSION

The number of users was incremented until page DASD utilization was driven to 75% on CDK-formatted DASD. Then, with the same number of users, the same workload was run on ECKD-formatted DASD.

When comparing ECKD-formatted paging DASD to CKD-formatted paging DASD, the following table shows that the performance results were equivalent. Furthermore, VMPRF data shows that the device utilization, service times, and response times for the paging DASD in the two different formats were equivalent.

PAGING DASD FORMAT RELEASE RUN ID	CKD ESA 1.1 Y34R1602	ECKD ESA 1.1 Y34R1603
Environment		
REAL STORAGE	**256M**	**256M**
EXP. STORAGE	**100M**	**100M**
USERS	**1600**	**1600**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**3**	**3**
Response Time		
TRIV INT	0.062	0.061
NONTRIV INT	0.322	0.314
TOT INT	0.225	0.218
TOT INT ADJ	0.234	0.227
AVG FIRST (T)	0.310	0.300
AVG LAST (T)	0.420	0.410
Throughput		
AVG THINK (T)	25.82	25.88
ETR	59.19	59.00
ETR (T)	56.94	56.74
ETR RATIO	1.039	1.040
ITR (H)	120.65	120.02
ITR	41.88	41.54
EMUL ITR	64.88	64.13
ITRR (H)	1.000	0.995
ITRR	1.000	0.992
Proc. Usage		
PBT/CMD (H)	24.866	24.995
PBT/CMD	24.762	25.028
CP/CMD (H)	9.125	9.110
CP/CMD	8.781	8.813
EMUL/CMD (H)	15.736	15.880
EMUL/CMD	15.981	16.215
Processor Util.		
TOTAL (H)	141.59	141.82
TOTAL	141.00	142.00
UTIL/PROC (H)	47.20	47.27
UTIL/PROC	47.00	47.33
TVR(H)	1.58	1.57
TVR	1.55	1.54
Storage		
WKSET (V)	74	75
PGBLPGS	56269	56270
PGBLPGS/USER	35.2	35.2
FREEPGS	4151	4139
FREE UTIL	0.95	0.95
SHRPGS	824	846
Paging		
READS/SEC	432	430
WRITES/SEC	269	271
PAGE/CMD	12.311	12.355
XSTOR IN/SEC	12	12
XSTOR OUT/SEC	257	267
XSTOR/CMD	4.724	4.917
FAST CLR/CMD	5.602	5.711

PAGING DASD FORMAT **RELEASE** **RUN ID**	**CKD** **ESA 1.1** **Y34R1602**	**ECKD** **ESA 1.1** **Y34R1603**
Environment		
REAL STORAGE	**256M**	**256M**
EXP. STORAGE	**100M**	**100M**
USERS	**1600**	**1600**
VTAMs	**1**	**1**
VSCSs	**2**	**2**
PROCESSORS	**3**	**3**
I/O		
VIO RATE	478	478
VIO/CMD	8.395	8.425
MDC READS	332	336
MDC WRITES	154	155
MDC MODS	128	128
MDC HIT RATIO	0.92	0.92
PRIVOPs		
PRIVOP/CMD	22.065	22.038
DIAG/CMD	23.639	23.966
DIAG 08/CMD	0.720	0.740
DIAG 10/CMD	0.000	0.000
DIAG 58/CMD	1.229	1.216
DIAG 98/CMD	0.281	0.282
DIAG A4/CMD	4.022	4.071
DIAG A8/CMD	1.862	1.815
DIAG 214/CMD	12.153	12.408
SIE/CMD	55.513	55.872
SIE INTCPT/CMD	41.080	41.346
FREE TOTL/CMD	83.770	85.994
VTAM Machines		
WKSET (V)	564	647
TOT CPU/CMD (V)	3.9058	3.9300
CP CPU/CMD (V)	1.9849	1.9954
VIRT CPU/CMD (V)	1.9209	1.9346
DIAG 98/CMD (V)	na	na
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 58. Comparing ECKD-Formatted DASD with CKD-Formmatted DASD

10. Tuning Considerations

Recommended 9221 Tuning

9221-170 / Minidisk

This section describes the performance benefits when tuning is applied to VM/ESA 1.1 for 9221 processors.

1) WORKLOAD: FS7B0R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9221-170
- STORAGE:
   - RSTOR:   48M,64M (see table)
   - XSTOR:   0M,16M  (see table), (all reserved for MDC)

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  H3AP01     3380
              H3SRV      3380
              H3RES      3380
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	3 - 3380-03	2	2	2	4	0
3380-D	1 - 3380-03	1	0	1	1	0

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINESPEED
3088-02	1	4.5M

3) SOFTWARE CONFIGURATION

```
- DRIVER:             TPNS
- THINK TIME DISTR:   BACTRIAN
- CMS BLOCKSIZE:      4K
- USER VM SIZE:       2M
- USER CMS MODE:      370
- USER RELSHARE:      100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAM	VTAM/VSCS	64M/XA	10000	QUICKDSP ON

4) MEASUREMENT DISCUSSION

The following table summarizes the results of tuning VM/ESA 1.1 measurements on the 9221-170 for the minidisk-only CMS intensive workloads.

When migrating from VM/ESA 1.0 370 Feature to VM/ESA 1.1 without tuning, a decrease in performance was experienced. For details, see section “9221-170 / Minidisk” on page 86. Consequently, tuning strategies were implemented for VM/ESA 1.1 which reduce DASD I/O, VTAM I/O, and SIE instructions. With this tuning, VM/ESA 1.1 internal throughput improved by 9.5% and external response time decreased by 38.2%. Also, CPU time for CP decreased by 14.1% and SIE/CMD improved by 19.0%.

Although specific measurements for each tuning option are not included in this report, an estimate of internal throughput improvement is provided.

- Minidisk Cache

  The untuned run had no expanded storage; therefore, it does not use minidisk caching. The tuned run had a portion of real storage configured as expanded storage and dedicated to minidisk caching. The tuned run showed 53 MDC READS from expanded storage replacing DASD I/Os, resulting in a decrease in CP/CMD(H). The estimated benefit to internal throughput was 2.1%.

  Comparing the 9221-170 to the 9121-480 (“Using XSTOR on a 9121” on page 228), the 9221-170 showed a better improvement in internal throughput and external response time.

  The command RETAIN XSTOR MDC ALL should be issued. The command RETAIN XSTOR MDC MIN MAX with MIN equal to MAX and MAX equal to 16M did not provide as much benefit. With the latter command, the system will use XSTOR frames not currently in use by MDC for paging. This results in page migration activity with very little benefit.

- DSPSLICE

  The default dispatch slice is too low for the 9221 because the mix of instructions sampled by CP′s timing loop (during CP initialization) is not representative of overall instruction execution time on this processor. Increasing DSPSLICE to three times the default reduces the number of timer interrupts for time slice end processing and reduces the associated SIE instructions. This tuning option contributed to decreasing CP/CMD (H) and SIE/CMD. The estimated benefit to internal throughput was 2.7%.

- VTAM Delay

  The VTAM delay was set to 0.2, resulting in a decrease in VTAM I/O which shows under VTAM Machines as a decrease in DIAG 98/CMD (V) of 16.2%. The VTAM I/O reduction caused a reduction in VTAM TOT CPU/CMD by 11.7%, VTAM VIRT CPU/CMD by 10.8%, and a reduction in unassisted privileged operations (PRIVOP/CMD). This tuning option contributed to decreasing CP processor usage per command (CP/CMD (H)) and SIE/CMD. External response time improved along with internal throughput. The estimated benefit to internal throughput was 3.2%.

- IPOLL ON

  IPOLL was turned on for VTAM causing IUCV instructions to be reduced. The result was a reduction in unassisted privileged operations

(PRIVOP/CMD), CP processor usage per command (CP/CMD (H)) and SIE/CMD. The estimated benefit to internal throughput was 0.2%

Turning IPOLL ON for VTAM showed similar results as recorded in section "GCS IPOLL Option" on page 207.

- Preloaded Saved Segments

  The FORTRAN and Script saved segments were preloaded before the measurement by issuing the SEGMENT LOAD command from an idle user during system startup. The 9221 did not have enough concurrent users using the saved segments. When no users are using the saved segments, the page frames become invalid and require a page read when the next user wants them. Preloading the saved segments in the tuned measurement accounted for most of the 25.3% decrease in READ/SEC. The decrease in READ/SEC occured even though the tuned case had less real storage. The estimated benefit to internal throughput was 1.3%.

9221 TUNING RELEASE RUN ID	NO ESA 1.1 H14R0283	YES ESA 1.1 H14R0287
Environment		
REAL STORAGE	**64M**	**48M**
EXP. STORAGE	**0M**	**16M**
USERS	**280**	**280**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
Response Time		
TRIV INT	0.226	0.172
NONTRIV INT	1.383	0.810
TOT INT	1.061	0.624
TOT INT ADJ	0.869	0.529
AVG FIRST (T)	0.320	0.270
AVG LAST (T)	0.890	0.550
Throughput		
AVG THINK (T)	28.41	28.39
ETR	7.89	8.23
ETR (T)	9.63	9.71
ETR RATIO	0.819	0.848
ITR (H)	10.86	11.90
ITR	8.91	10.10
EMUL ITR	14.48	15.99
ITRR (H)	1.000	1.095
ITRR	1.000	1.134
Proc. Usage		
PBT/CMD (H)	92.057	84.060
PBT/CMD	92.420	84.486
CP/CMD (H)	42.339	36.380
CP/CMD	35.306	30.910
EMUL/CMD (H)	49.718	47.680
EMUL/CMD	57.113	53.577
Processor Util.		
TOTAL (H)	88.65	81.59
TOTAL	89.00	82.00
UTIL/PROC (H)	88.65	81.59
UTIL/PROC	89.00	82.00
TVR(H)	1.85	1.76
TVR	1.62	1.58
Storage		
WKSET (V)	76	77
PGBLPGS	13585	9520
PGBLPGS/USER	48.5	34.0
FREEPGS	804	804
FREE UTIL	0.88	0.88
SHRPGS	785	900
Paging		
READS/SEC	83	62
WRITES/SEC	48	49
PAGE/CMD	13.603	11.437
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	5.711	5.461

9221 TUNING RELEASE RUN ID	NO ESA 1.1 H14R0283	YES ESA 1.1 H14R0287
Environment		
REAL STORAGE	**64M**	**48M**
EXP. STORAGE	**0M**	**16M**
USERS	**280**	**280**
VTAMs	**1**	**1**
VSCSs	**0**	**0**
PROCESSORS	**1**	**1**
I/O		
VIO RATE	80	82
VIO/CMD	8.307	8.449
MDC READS	0	53
MDC WRITES	0	26
MDC MODS	0	21
MDC HIT RATIO	0.00	0.91
PRIVOPs		
PRIVOP/CMD	18.028	14.426
DIAG/CMD	26.436	26.002
DIAG 08/CMD	0.623	0.618
DIAG 10/CMD	0.000	0.000
DIAG 58/CMD	1.142	1.236
DIAG 98/CMD	2.908	2.473
DIAG A4/CMD	3.946	3.812
DIAG A8/CMD	1.765	2.061
DIAG 214/CMD	11.423	11.334
SIE/CMD	70.301	56.977
SIE INTCPT/CMD	47.102	42.163
FREE TOTL/CMD	101.350	97.571
VTAM Machines		
WKSET (V)	207	211
TOT CPU/CMD (V)	20.7685	18.3362
CP CPU/CMD (V)	9.3869	8.1785
VIRT CPU/CMD (V)	11.3816	10.1577
DIAG 98/CMD (V)	3.004	2.519
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM		

Table 59. 9221-170 Tuning

Using XSTOR on a 9121

9121-480 / 35% SFS

This section describes the performance tradeoffs when using a portion of real storage for expanded storage used exclusively for minidisk caching.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

- PROCESSOR: 9121-480
- STORAGE:
  - RSTOR: 192M,256M (see table)
  - XSTOR: 0M,64M (see table), all reserved for MDC

- DASD:

	PACK NAME	TYPE
- SYSTEM:	PSYS02	3380-A
	PSPT01	3380-A
	WKLD01	3380-A
	WKLD02	3380-A

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-A	11 - 3880-2	16	4	8	0	20
3380-A	2 - 3880-J23	0	0	0	4	0

- TAPE: MONITOR 3480

- COMMUNICATIONS:

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	2	44	56Kb

3) SOFTWARE CONFIGURATION

- DRIVER: TPNS
- THINK TIME DISTR: BACTRIAN
- CMS BLOCKSIZE: 4K
- USER VM SIZE: 2M
- USER CMS MODE: XA
- USER RELSHARE: 100

- SERVER MACHINES:

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON, RESERVE 850
RWSERV1	SFS	32M/XA	1500	QUICKDSP ON, RESERVE 1300
RWSERV2	SFS	32M/XA	1500	QUICKDSP ON, RESERVE 1300
CRRSERV1	CRR	17M/XA	100	QUICKDSP ON

4) MEASUREMENT DISCUSSION

The first run in the following table had no expanded storage; thus, it did not use minidisk caching and all paging was to DASD. The second run shown here had a portion of real storage used for expanded storage which in turn, was used exclusively for minidisk caching. These two runs show very similar performance characteristics with the external response time (AVE LAST (T)) of the minidisk caching run being slightly better. As expected, the use of minidisk caching improved I/O time but caused the paging rate to increase. This is shown in the table as a slightly worse paging rate per command (PAGE/CMD) for the minidisk caching run and an improved SFS server I/O time per command (I/O TIME/CMD (Q)).

Thus, for the FS7B workload, the use of a portion of real storage for minidisk caching provided slightly better performance. However, this may not be true for every installation. FS7B is a very uniform workload which provides for a high minidisk cache hit ratio and a relatively low demand on real storage. This translates into smaller real storage requirements and better response times due to the replacement of minidisk I/O with faster page I/O. Installations with expanded storage available (this machine had no true expanded storage available) should benefit from the use of some of that storage for minidisk caching. If real storage is tight, installations may want to run without using a portion of real storage as expanded storage for minidisk caching to reduce paging, especially if their workload doesn't provide for a high minidisk cache hit ratio.

EXP. STORAGE USE RELEASE RUN ID	NONE ESA 1.0 L23F1482	MDC ONLY ESA 1.0 L23F1481
Environment		
REAL STORAGE	**256M**	**192M**
EXP. STORAGE	**0M**	**64M**
USERS	**1480**	**1480**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**2**	**2**
Response Time		
TRIV INT	0.074	0.072
NONTRIV INT	0.901	0.867
TOT INT	0.581	0.552
TOT INT ADJ	0.633	0.609
AVG FIRST (T)	0.425	0.445
AVG LAST (T)	0.875	0.870
Throughput		
AVG THINK (T)	25.48	25.42
ETR	57.13	57.92
ETR (T)	52.40	52.50
ETR RATIO	1.090	1.103
ITR (H)	57.92	58.65
ITR	31.63	32.28
EMUL ITR	52.67	53.81
ITRR (H)	1.000	1.012
ITRR	1.000	1.020
Proc. Usage		
PBT/CMD (H)	34.528	34.103
PBT/CMD	34.544	34.289
CP/CMD (H)	14.050	13.930
CP/CMD	13.741	13.716
EMUL/CMD (H)	20.470	20.164
EMUL/CMD	20.803	20.573
Processor Util.		
TOTAL (H)	180.92	179.03
TOTAL	181.00	180.00
UTIL/PROC (H)	90.46	89.51
UTIL/PROC	90.50	90.00
TVR(H)	1.69	1.69
TVR	1.66	1.67
Storage		
WKSET (V)	75	73
PGBLPGS	57590	41111
PGBLPGS/USER	38.9	27.8
FREEPGS	3742	3746
FREE UTIL	0.94	0.94
SHRPGS	1120	1096
Paging		
READS/SEC	396	421
WRITES/SEC	297	313
PAGE/CMD	13.226	13.982
XSTOR IN/SEC	0	0
XSTOR OUT/SEC	0	0
XSTOR/CMD	0.000	0.000
FAST CLR/CMD	7.863	7.772

EXP. STORAGE USE RELEASE RUN ID	NONE ESA 1.0 L23F1482	MDC ONLY ESA 1.0 L23F1481
Environment		
REAL STORAGE	**256M**	**192M**
EXP. STORAGE	**0M**	**64M**
USERS	**1480**	**1480**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**2**	**2**
I/O		
VIO RATE	356	360
VIO/CMD	6.794	6.858
MDC READS	0	283
MDC WRITES	0	100
MDC MODS	0	60
MDC HIT RATIO	0.00	0.87
PRIVOPs		
PRIVOP/CMD	29.453	29.083
DIAG/CMD	15.806	15.831
DIAG 08/CMD	0.744	0.743
DIAG 10/CMD	5.210	5.200
DIAG 58/CMD	1.241	1.238
DIAG 98/CMD	0.534	0.514
DIAG A4/CMD	2.538	2.572
DIAG A8/CMD	1.737	1.753
DIAG 214/CMD	na	na
SIE/CMD	79.853	77.207
SIE INTCPT/CMD	55.897	55.589
FREE TOTL/CMD	140.468	133.803
VTAM Machines		
WKSET (V)	1060	1011
TOT CPU/CMD (V)	5.2591	5.1751
CP CPU/CMD (V)	2.9900	2.9103
VIRT CPU/CMD (V)	2.2690	2.2648
DIAG 98/CMD (V)	0.535	0.533
SFS Servers		
WKSET (V)	1367	1362
TOT CPU/CMD (V)	4.4638	4.4660
CP CPU/CMD (V)	2.2584	2.2965
VIRT CPU/CMD (V)	2.2054	2.1695
FP REQ/CMD(Q)	1.333	1.338
IO/CMD (Q)	1.855	1.819
IO TIME/CMD (Q)	0.075	0.068
SFS TIME/CMD (Q)	0.210	0.213
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Q=Filepool Counters, Unmarked=RTM		

Table 60. 9121-480 Tuning With a Portion of Real Storage Used as Expanded Storage Exclusively for Minidisk Caching.

Set Reserve Option

9121-480 / 35% SFS

This section describes the performance gain when using the SET RESERVE option for VM servers on the 9121-480 with limited real storage.

1) WORKLOAD: FS7B35R

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9121-480
- STORAGE:
   - RSTOR:       192M
   - XSTOR:        64M     All reserved for MDC

- DASD:
            PACK NAME   TYPE
   - SYSTEM:  PSYS02      3380-A
              PSPT01      3380-A
              WKLD01      3380-A
              WKLD02      3380-A
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	PAGE	SPOOL	TDISK	USER	SERVER
3380-A	11 - 3880-2	16	4	8	0	20
3380-A	2 - 3880-J23	0	0	0	0	4

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	2	44	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR:  BACTRIAN
- CMS BLOCKSIZE:     4K
- USER VM SIZE:      2M
- USER CMS MODE:     XA
- USER RELSHARE:     100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON, RESERVE OFF/850
RWSERV1	SFS	32M/XA	1500	QUICKDSP ON, RESERVE OFF/1300
RWSERV2	SFS	32M/XA	1500	QUICKDSP ON, RESERVE OFF/1300
CRRSERV1	CRR	17M/XA	100	NONE

4) MEASUREMENT DISCUSSION

The first run in the table had only the RELSHARE and QUICKDSP ON tuning options set for the servers, while the second run also had the servers′ working set reserved in storage with the SET RESERVE command. The number of pages reserved for each server is shown in the preceding table. As indicated by an external response time (AVG LAST (T)) over 36 seconds, the performance of the first run was unacceptable. Further inspection of the VMPRF data for the run indicated a serial page fault problem for the SFS and (to a lesser degree) the VSCS servers. The applicable VMPRF data for both runs is shown in the following table:

	Percent of True Non-Dormant Time						Page Reads + Writes/Sec	
	Running		Page Wait		Utilization			
Run #	**1**	**2**	**1**	**2**	**1**	**2**	**1**	**2**
SFS	1.9	4.9	63.8	15.6	66.7	20.5	75.2	38.7
VSCS	3.8	9.7	29.0	16.4	32.8	26.1	31.2	29.4
Note: Utilization=Running+Page Wait								

Table 61. SFS and VSCS Server States.

This table shows that over 63% of the non-dormant time the SFS servers were in page wait; therefore, the users were waiting for the servers′ pages to be brought into storage which, in effect, serializes the server and all its dependent users and increases response times.

This environment was susceptible to this serial page fault phenomenon since it had little or no expanded storage for paging and all page reads and writes were to DASD. After determining that serial page faulting was occurring in the first run, it was decided to reserve the SFS and VSCS pages in storage.

With this change, the percent of time in page wait for the SFS and VSCS servers improved dramatically and the paging rate to DASD for the servers was cut almost in half. External response times are 0.945, much improved over the original run. The paging rate (PAGE/CMD) was reduced by two-thirds and the amount of time per command spent in the SFS server (SFS TIME/CMD) was greatly improved.

There is a potential downside to using the SET RESERVE command not shown here. Reserving pages for a given user may cause other users to experience increased paging due to fewer pages left in the Dynamic Paging Area (DPA). Care must be taken not to reserve more pages than is needed by the virtual machine. SET RESERVE should most often be used for those virtual machines that, when taking a page fault, will degrade the performance of more than just that particular virtual machine. Examples include servers and guest operating systems.

SET RESERVE RELEASE RUN ID	OFF ESA 1.0 L23F1480	850/1300 ESA 1.0 L23F1484
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**1480**	**1480**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**2**	**2**
Response Time		
TRIV INT	0.162	0.071
NONTRIV INT	69.251	0.895
TOT INT	38.821	0.564
TOT INT ADJ	48.404	0.636
AVG FIRST (T)	9.500	0.495
AVG LAST (T)	36.675	0.945
Throughput		
AVG THINK (T)	26.26	25.52
ETR	23.92	58.77
ETR (T)	19.18	52.10
ETR RATIO	1.247	1.128
ITR (H)	43.77	58.55
ITR	28.59	33.01
EMUL ITR	58.51	55.46
ITRR (H)	1.000	1.338
ITRR	1.000	1.155
Proc. Usage		
PBT/CMD (H)	45.695	34.158
PBT/CMD	43.786	34.168
CP/CMD (H)	23.668	14.083
CP/CMD	22.414	13.821
EMUL/CMD (H)	22.004	20.068
EMUL/CMD	21.372	20.347
Processor Util.		
TOTAL (H)	87.66	177.95
TOTAL	84.00	178.00
UTIL/PROC (H)	43.83	88.97
UTIL/PROC	42.00	89.00
TVR(H)	2.08	1.70
TVR	2.05	1.68
Storage		
WKSET (V)	56	73
PGBLPGS	41496	41095
PGBLPGS/USER	28.0	27.8
FREEPGS	3624	3748
FREE UTIL	0.94	0.94
SHRPGS	1142	1067
Paging		
READS/SEC	432	403
WRITES/SEC	348	318
PAGE/CMD	40.658	13.840
XSTOR IN/SEC	41	0
XSTOR OUT/SEC	99	0
XSTOR/CMD	7.298	0.000
FAST CLR/CMD	7.662	7.832

SET RESERVE	**OFF**	**850/1300**
RELEASE	**ESA 1.0**	**ESA 1.0**
RUN ID	**L23F1480**	**L23F1484**
Environment		
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
USERS	**1480**	**1480**
VTAMs	**1**	**1**
VSCSs	**1**	**1**
PROCESSORS	**2**	**2**
I/O		
VIO RATE	126	357
VIO/CMD	6.568	6.853
MDC READS	115	278
MDC WRITES	44	96
MDC MODS	29	56
MDC HIT RATIO	0.88	0.87
PRIVOPs		
PRIVOP/CMD	31.182	28.974
DIAG/CMD	17.409	15.849
DIAG 08/CMD	0.678	0.749
DIAG 10/CMD	5.004	5.183
DIAG 58/CMD	1.199	1.248
DIAG 98/CMD	0.990	0.518
DIAG A4/CMD	2.398	2.495
DIAG A8/CMD	1.668	1.843
DIAG 214/CMD	na	na
SIE/CMD	85.643	77.626
SIE INTCPT/CMD	58.237	55.115
FREE TOTL/CMD	168.157	133.485
VTAM Machines		
WKSET (V)	863	1072
TOT CPU/CMD (V)	5.8529	5.1934
CP CPU/CMD (V)	3.5060	2.9646
VIRT CPU/CMD (V)	2.3470	2.2288
DIAG 98/CMD (V)	1.040	0.536
SFS Servers		
WKSET (V)	1194	1364
TOT CPU/CMD (V)	5.5921	4.4256
CP CPU/CMD (V)	3.0713	2.2715
VIRT CPU/CMD (V)	2.5208	2.1542
FP REQ/CMD(Q)	1.219	1.328
IO/CMD (Q)	2.351	1.825
IO TIME/CMD (Q)	2.273	0.071
SFS TIME/CMD (Q)	4.618	0.224

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

Table 62. Tuning on 9121-480 with SET RESERVE.

OfficeVision MSGFLAGS Settings

This section explores the performance benefits of using the existing tuning option within OfficeVision of turning off the console messages for the Calendar and Mailbox service machines. This ability is documented in *Managing OfficeVision/VM* under the topic of MSGFLAGS.

9021-720

The following 9021-720 runs are provided to show the benefits of tuning OfficeVision by setting messages OFF for the Calendar and Mailbox servers.

1) WORKLOAD: IOB V2.1

2) HARDWARE CONFIGURATION

```
- PROCESSOR:  9021-720
- STORAGE:
   - RSTOR:   512 M
   - XSTOR:   2048 M

- DASD:
              PACK NAME  TYPE
   - SYSTEM:  PSYS02     3380-A
              PSPT01     3380-D
              WKLD01     3380-D
              WKLD02     3380-D
```

TYPE OF DASD	NUMBER/TYPE OF CONTROL UNIT	NUMBER OF PACKS PAGE	SPOOL	TDISK	USER	SERVER
3380-D	20 - 3880-3	20	16	12	40	0
3380-D	1 - 3880-G23	0	0	0	0	4
3380-A	3 - 3880-G23	0	0	0	0	12
3380-K	2 - 3990-2	0	0	0	0	16

```
- TAPE:       MONITOR  3480

- COMMUNICATIONS:
```

CONTROLLER	NUMBER	LINES/CONTROLLER	LINESPEED
3745-410	3	36	56Kb

3) SOFTWARE CONFIGURATION

```
- DRIVER:            TPNS
- THINK TIME DISTR:  IOB
- CMS BLOCKSIZE:     4K
- USER VM SIZE:      2M
- USER CMS MODE:     XA
- USER RELSHARE:     100

- SERVER MACHINES:
```

SERVER MACHINE	TYPE	VM SIZE/ CMS MODE	RELSHARE	OTHER OPTIONS
VTAMXAA	VTAM	64M/XA	10000	QUICKDSP ON; IPOLL ON
VSCSXA2	VSCS	64M/XA	10000	QUICKDSP ON; IPOLL ON
VSCSXA3	VSCS	64M/XA	10000	QUICKDSP ON; IPOLL ON
PRODBM	OV/VM	16M/XA	10000	QUICKDSP ON
PROCAL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMAIL	OV/VM	16M/XA	10000	QUICKDSP ON
PROMBX00 - 50	OV/VM	16M/XA	10000	QUICKDSP ON; IBCENTRL=Y

4) MEASUREMENT DISCUSSION

The following is a brief description of the measurements discussed in this section. In all three runs, the VTAM and VSCS servers had “SET IPOLL ON” issued before they were set up. CAL/MBX MSG OFF indicates that console messages were turned off by issuing the “MSGFLAGS FF” command for the Calendar server machine and the “MSGFLAGS FFFF” command for the Mailbox server machines.

Y64V620B 6200 users on VM/ESA 1.1 with CAL/MBX MSG ON

Y64V620C 6200 users on VM/ESA 1.1 with CAL/MBX MSG OFF

Y64V6401 6400 users on VM/ESA 1.1 with CAL/MBX MSG OFF

Prior measurements on a 6-way processor showed some indications of a master processor bottleneck. This was observed by looking at the amount of master processor emulation CPU time that is consumed as compared to the other processors in the configuration. As more master processor CP CPU is required, the amount of master emulation that can take place is reduced. One factor contributing to master processor usage is console I/O activity.

A set of measurements were made to see what effect these messages have on system capacity in this 6-way environment. First, a comparison between messages on (default) and messages off was performed at the same number of users. The results showed a reduction of 0.06 seconds (7%) in external response time (AVG LAST (T)) and a 3% reduction in CPU utilization (TOTAL (H)) for this workload. The percent of emulation time on the master processor went from 6.3% to 9.9% (this data was extracted from the VMPRF reports and is not included in the data charts below). This increase in emulation utilization implies that more user work is being allowed to run on the master processor.

Second, a measurement was made with the messages turned off and the number of users increased to 6400. The percent of emulation utilization on the master processor decreased from 9.9% to 8.9% with the addition of the extra users. The CPU utilization was approximately equivalent to the measurement with messages turned on, yet master emulation was still improved (6.3% to 8.9%).

Therefore, it is concluded that turning messages off allowed for improved capacity due to reduced resources consumed and potentially reduced master processor requirements.

MSGFLAGS SETTING RELEASE RUN ID	MSG ON ESA 1.1 Y64V620B	MSG OFF ESA 1.1 Y64V620C	MSG OFF ESA 1.1 Y64V6401
Environment			
REAL STORAGE	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**
USERS	**6200**	**6200**	**6400**
VTAMs	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**
Response Time			
TRIV INT	0.036	0.032	0.035
NONTRIV INT	0.401	0.349	0.385
TOT INT	0.302	0.266	0.290
TOT INT ADJ	0.371	0.322	0.357
AVG FIRST (T)	0.627	0.580	0.670
AVG LAST (T)	0.833	0.773	0.910
Throughput			
AVG THINK (T)	42.37	42.55	42.19
ETR	141.91	139.77	146.59
ETR (T)	115.38	115.36	118.93
ETR RATIO	1.230	1.212	1.233
ITR (H)	129.69	133.66	134.31
ITR	26.62	27.04	27.67
EMUL ITR	48.79	49.07	50.02
ITRR (H)	1.000	1.031	1.036
ITRR	1.000	1.016	1.040
Proc. Usage			
PBT/CMD (H)	46.264	44.889	44.673
PBT/CMD	46.282	44.902	44.564
CP/CMD (H)	22.514	21.549	21.422
CP/CMD	21.061	20.197	19.927
EMUL/CMD (H)	23.746	23.334	23.246
EMUL/CMD	25.221	24.705	24.636
Processor Util.			
TOTAL (H)	533.80	517.85	531.31
TOTAL	534.00	518.00	530.00
UTIL/PROC (H)	88.97	86.31	88.55
UTIL/PROC	89.00	86.33	88.33
TVR(H)	1.95	1.92	1.92
TVR	1.84	1.82	1.81
Storage			
WKSET (V)	53	53	53
PGBLPGS	108K	108K	108K
PGBLPGS/USER	17.8	17.8	17.3
FREEPGS	13014	13034	13520
FREE UTIL	0.96	0.96	0.97
SHRPGS	1411	1394	1350
Paging			
READS/SEC	30	34	51
WRITES/SEC	74	78	84
PAGE/CMD	0.901	0.971	1.135
XSTOR IN/SEC	1659	1603	1687
XSTOR OUT/SEC	1813	1765	1865
XSTOR/CMD	30.092	29.195	29.866
FAST CLR/CMD	14.977	14.988	14.874

MSGFLAGS SETTING RELEASE RUN ID	MSG ON ESA 1.1 Y64V620B	MSG OFF ESA 1.1 Y64V620C	MSG OFF ESA 1.1 Y64V6401
Environment			
REAL STORAGE	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**
USERS	**6200**	**6200**	**6400**
VTAMs	**1**	**1**	**1**
VSCSs	**2**	**2**	**2**
PROCESSORS	**6**	**6**	**6**
I/O			
VIO RATE	2514	2385	2440
VIO/CMD	21.789	20.674	20.516
MDC READS	1379	1375	1410
MDC WRITES	987	976	995
MDC MODS	841	830	844
MDC HIT RATIO	0.90	0.90	0.90
PRIVOPs			
PRIVOP/CMD	18.145	17.422	17.125
DIAG/CMD	85.526	83.875	83.286
DIAG 08/CMD	9.430	9.197	8.955
DIAG 10/CMD	0.000	0.000	0.000
DIAG 58/CMD	1.863	1.864	1.867
DIAG 98/CMD	0.416	0.407	0.404
DIAG A4/CMD	11.666	11.607	11.528
DIAG A8/CMD	6.275	6.207	6.138
DIAG 214/CMD	34.989	34.873	34.886
SIE/CMD	124.250	124.270	120.540
SIE INTCPT/CMD	79.520	79.533	75.940
FREE TOTL/CMD	230.750	221.911	215.250
VTAM Machines			
WKSET (V)	1788	1802	1847
TOT CPU/CMD (V)	5.3371	5.1789	5.0449
CP CPU/CMD (V)	2.2541	2.3103	2.2655
VIRT CPU/CMD (V)	3.0830	2.8686	2.7794
DIAG 98/CMD (V)	0.418	0.415	0.407
Note: T=TPNS, V=VMPRF, H=Hardware Monitor, Unmarked=RTM			

Table 63. The Effects of the OfficeVision MSGFLAGS Setting

Part 4. Appendixes

Appendix A. CMS Trace Data

Measurement Methodology

A selected set of CMS commands were traced with VM/ESA 1.0 and VM/ESA 1.1 to compare their virtual pathlengths, shared and non-shared page usage, and counts of special operations executed. Special operations are those instructions which cause either CP or the hardware microcode to provide services for the virtual machine. These include privileged instructions, DIAGNOSE instructions, and the first time a non-shared page is referenced. Three different environments were traced: CMS with minidisks, CMS with SFS, and CMS with SFS using a VM Data Space. In the minidisk and SFS environments, traces were done for 370, XA, and (for VM/ESA 1.1) XC mode virtual machines.

The technique used to collect the trace data involves issuing the selected commands from a 3270 terminal while the CP TRACE command 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 CP TRACE command. 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.

Note: The TRREAD and STARS programs mentioned above are internal tools and can not be ordered by customers.

Minidisk (EDF)

Commands Traced

A set of twenty-five CMS commands were traced for both VM/ESA 1.0 and VM/ESA 1.1. 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, minidisk 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 commands that were traced:

```
SET       IMSG OFF
ACCESS    295 C
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
 FILE
ERASE     TEST SCRIPT A
EXEC2     X Y (ALL
QUERY     DISK
FILEDEF   IN READER
PRINT     MASTER SCRIPT A
ASSEMBLE  BR14
LOAD      BR14
GENMOD    BR14
BR14
REXREX    X Y (ALL
SPKA
```

Results

Pathlength

RELEASE	1.0		1.1		
USER MODE	370	XA	370	XA	XC
SET	9225	10282	9404	10454	10503
ACCESS	17940	19777	18148	19907	20018
COPY	28598	29753	28610	29902	30041
COMPARE	66756	69595	64213	66970	67131
RENAME	20879	22043	21172	22743	22805
LISTFILE	10406	12221	10687	12132	12176
XEDIT	76452	78093	75524	77016	77147
XXXX	30198	31137	30665	31527	31614
LOCATE	49347	50008	49468	50105	50143
CHANGE	20241	20902	20574	20985	21028
NEXT	37068	37729	37189	37838	37876
INPUT	101210	102586	102244	102873	102974
DELETE	27190	27851	27311	27902	27940
FILE	63792	65008	62630	63623	63758
ERASE	15451	16443	15685	16880	16940
EXEC2	64115	66586	64332	66683	67134
QUERY	19921	24628	20348	24585	24695
FILEDEF	6833	7981	6954	7998	8026
PRINT	50448	52014	49310	50943	51024
ASSEMBLE	360004	371529	364607	374318	377076
LOAD	16925	18033	20873	21617	21834
GENMOD	14535	15635	14841	15722	15849
BR14	12784	14052	12996	14164	14292
REXREX	68984	70646	70932	72656	72838
SPKA	59935	63492	60083	63094	63402
TOTAL	1249237	1298024	1258800	1302637	1308264
AVERAGE	49969	51921	50352	52105	52331

Non-Shared Pages

RELEASE	1.0		1.1		
USER MODE	370	XA	370	XA	XC
SET	23	23	23	23	23
ACCESS	22	22	23	22	22
COPY	30	30	31	32	32
COMPARE	24	24	24	24	24
RENAME	28	28	28	28	28
LISTFILE	24	24	27	26	26
XEDIT	44	44	46	46	46
XXXX	25	26	26	27	27
LOCATE	20	21	22	24	24
CHANGE	20	20	23	25	25
NEXT	19	20	21	23	23
INPUT	24	24	27	29	29
DELETE	19	20	21	24	24
FILE	36	35	36	38	38
ERASE	28	28	28	28	28
EXEC2	18	18	18	18	18
QUERY	23	23	23	23	23
FILEDEF	22	23	22	23	23
PRINT	27	27	27	27	27
ASSEMBLE	102	102	99	101	101
LOAD	38	38	41	42	42
GENMOD	35	35	35	35	35
BR14	24	24	24	25	25
REXREX	21	21	21	21	21
SPKA	38	38	39	39	39
TOTAL	734	738	755	773	773
AVERAGE	29.4	29.5	30.2	30.9	30.9

Shared Pages

RELEASE	1.0		1.1		
USER MODE	370	XA	370	XA	XC
SET	42	45	41	44	44
ACCESS	57	59	55	60	60
COPY	49	51	49	51	51
COMPARE	42	44	42	44	44
RENAME	48	50	48	50	50
LISTFILE	39	41	39	42	42
XEDIT	80	82	80	82	82
XXXX	67	71	67	71	71
LOCATE	41	44	42	44	44
CHANGE	51	54	52	55	55
NEXT	39	42	40	42	42
INPUT	59	62	58	63	63
DELETE	42	45	43	45	45
FILE	72	75	74	77	77
ERASE	48	49	47	48	48
EXEC2	37	38	36	37	37
QUERY	48	50	49	51	51
FILEDEF	32	34	31	33	33
PRINT	55	57	54	56	56
ASSEMBLE	71	73	74	77	77
LOAD	53	54	56	57	57
GENMOD	47	48	47	48	48
BR14	47	49	47	49	49
REXREX	45	46	48	49	49
SPKA	60	62	63	65	65
TOTAL	1271	1325	1282	1340	1340
AVERAGE	50.8	53.0	51.3	53.6	53.6

Special Operations

RELEASE	1.0		1.1		
USER MODE	370	XA	370	XA	XC
SET	45	88	45	80	91
ACCESS	79	160	73	124	150
COPY	84	129	76	113	156
COMPARE	254	589	252	570	613
RENAME	57	99	54	87	100
LISTFILE	53	119	57	106	115
XEDIT	107	203	106	157	185
XXXX	65	125	66	98	114
LOCATE	48	92	50	77	81
CHANGE	49	92	53	79	85
NEXT	47	91	49	76	80
INPUT	88	179	96	149	167
DELETE	47	91	49	77	81
FILE	114	182	99	140	173
ERASE	57	99	54	87	100
EXEC2	293	715	293	701	874
QUERY	100	278	100	239	256
FILEDEF	43	86	43	78	84
PRINT	63	110	59	94	115
ASSEMBLE	1023	1864	946	1236	1977
LOAD	85	128	88	129	190
GENMOD	78	125	72	107	146
BR14	58	100	57	91	132
REXREX	109	246	111	234	299
SPKA	184	506	186	352	388
TOTAL	3230	6496	3134	5281	6752
AVERAGE	129	260	125	211	270

Shared File System (SFS)

Commands Traced

For SFS, CMS commands are traced that would typically be issued by a user of SFS file pools. They are similar in function to the above commands used with minidisk. In addition, several commands which are unique to SFS are included. SFS unique commands that are measured are CREATE ALIAS, GRANT AUTHORITY, and REVOKE AUTHORITY. The SFS trace results are broken down by the virtual machine to which the processing 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   RWSERV1:OPERATOR.F100  H/A
COPY     A100A ASSEMBLE D MASTER FILE A
RENAME   MASTER FILE A TEST ASSEMBLE A
XEDIT    TEST ASSEMBLE A
 FILE
ERASE    TEST ASSEMBLE 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
```

These commands appear in subsequent tables prefixed either with U to denote user virtual machine activity or S to denote server virtual machine activity.

SFS User Results

Pathlength

RELEASE	1.0		1.1		
SERVER MODE	XA		XC		
USER MODE	370	XA	370	XA	XC
UACCESS	33738	35412	35058	36783	36880
UCOPY	19510	20724	22758	23848	23917
URENAME	22568	23717	25785	27303	27359
UXEDIT	68392	69728	70529	71746	71922
UFILE	71806	72780	73801	74773	74961
UERASE	17926	18930	21090	22279	22341
UASSEMBL	389444	401098	410978	421050	424001
ULOAD	29659	30763	32454	33490	33686
UGENMOD	21329	22491	22742	23745	23880
UBR14	21103	22442	22353	23569	23715
UCREATE	18315	19292	21590	22712	22765
UGRANT	37369	39200	36191	38301	38416
UREVOKE	37662	39494	36673	38785	38898
TOTAL	788821	816071	832002	858384	862741
AVERAGE	60679	62775	64000	66030	66365

Non-Shared Pages

RELEASE	1.0		1.1		
SERVER MODE	XA		XC		
USER MODE	370	XA	370	XA	XC
UACCESS	30	30	34	33	33
UCOPY	29	29	29	29	29
URENAME	31	31	32	32	32
UXEDIT	58	59	63	62	62
UFILE	60	60	63	63	63
UERASE	29	29	31	32	32
UASSEMBL	113	108	116	114	114
ULOAD	39	39	44	44	44
UGENMOD	34	34	36	36	36
UBR14	30	31	33	33	33
UCREATE	28	28	31	31	31
UGRANT	42	43	44	44	44
UREVOKE	36	37	38	38	38
TOTAL	559	558	594	591	591
AVERAGE	43.0	42.9	45.7	45.5	45.5

Shared Pages

RELEASE	1.0		1.1		
SERVER MODE	XA		XC		
USER MODE	370	XA	370	XA	XC
UACCESS	80	82	88	93	93
UCOPY	63	65	78	80	80
URENAME	68	70	80	82	82
UXEDIT	114	116	123	125	125
UFILE	112	115	121	124	124
UERASE	72	73	85	86	86
UASSEMBL	120	122	128	130	130
ULOAD	86	87	93	94	94
UGENMOD	81	82	86	87	87
UBR14	79	81	83	85	85
UCREATE	73	74	85	86	86
UGRANT	94	96	99	101	101
UREVOKE	93	95	99	101	101
TOTAL	1135	1158	1248	1274	1274
AVERAGE	87.3	89.1	96.0	98.0	98.0

Special Operations

RELEASE	1.0		1.1		
SERVER MODE	XA		XC		
USER MODE	370	XA	370	XA	XC
UACCESS	76	145	80	123	146
UCOPY	56	101	63	99	119
URENAME	55	97	62	95	109
UXEDIT	152	249	147	197	246
UFILE	149	218	142	181	239
UERASE	53	95	61	95	113
UASSEMBL	1055	1856	1006	1279	2106
ULOAD	97	140	99	133	200
UGENMOD	82	129	81	116	157
UBR14	80	123	82	115	166
UCREATE	52	94	61	94	106
UGRANT	93	181	95	144	170
UREVOKE	87	175	89	138	164
TOTAL	2087	3603	2068	2809	4041
AVERAGE	161	277	159	216	311

SFS Server Results

Pathlength

RELEASE	1.0		1.1		
SERVER MODE	XA		XC		
USER MODE	370	XA	370	XA	XC
SACCESS	318354	318354	332916	332916	332915
SCOPY	327870	311356	346202	348488	321623
SRENAME	39113	39113	39976	39974	39958
SXEDIT	90213	90235	95374	95368	94562
SFILE	603571	571817	382126	379611	351775
SERASE	73325	73334	74903	74897	74945
SASSEMBL	543318	543361	553651	551189	552117
SLOAD	70959	70961	73157	73123	72657
SGENMOD	58488	58488	59628	59606	59670
SBR14	16029	16032	16501	16499	16584
SCREATE	75277	75276	83657	83620	83626
SGRANT	62811	62812	64058	64043	64044
SREVOKE	63954	63954	69267	69251	69253
TOTAL	2343282	2295093	2191416	2188585	2133729
AVERAGE	180252	176546	168570	168353	164133

Sum of shared and non-shared pages

RELEASE	1.0		1.1		
SERVER MODE	XA		XC		
USER MODE	370	XA	370	XA	XC
SACCESS	114	114	121	121	121
SCOPY	184	183	193	192	192
SRENAME	106	106	117	117	117
SXEDIT	177	177	185	186	186
SFILE	207	206	221	220	219
SERASE	150	150	161	160	160
SASSEMBL	173	173	182	182	182
SLOAD	132	132	141	141	141
SGENMOD	125	125	138	138	138
SBR14	89	89	97	96	96
SCREATE	122	122	136	136	136
SGRANT	131	131	141	142	142
SREVOKE	128	128	139	139	139
TOTAL	1838	1836	1972	1970	1969
AVERAGE	141.4	141.2	151.7	151.5	151.5

Special Operations

RELEASE	1.0		1.1		
SERVER MODE	XA		XC		
USER MODE	370	XA	370	XA	XC
SACCESS	54	54	57	57	57
SCOPY	245	245	293	344	355
SRENAME	78	78	82	82	82
SXEDIT	350	350	356	356	356
SFILE	508	520	479	479	479
SERASE	118	118	130	130	130
SASSEMBL	2084	2084	2059	1963	1939
SLOAD	238	238	247	247	223
SGENMOD	140	140	146	146	146
SBR14	91	91	71	71	71
SCREATE	78	78	132	132	132
SGRANT	78	78	82	82	82
SREVOKE	78	78	82	82	82
TOTAL	4140	4152	4216	4171	4134
AVERAGE	318	319	324	321	318

VM Data Spaces

Commands Traced

For SFS data in VM Data Spaces, the CMS ACCESS command and XEDIT of a file are traced. There are two environments that are of interest:

1. data in a DIRCONTROL directory in a VM data space;
2. data in a FILECONTROL directory.

The directories have 500 equal size files. The commands will show virtual machine activity in both user and server machines. The trace data is broken down by the virtual machine to which the overhead is charged. The trace of the ACCESS command of a DIRCONTROL directory in a data space is the first ACCESS of that directory by that user. However, that directory has already been ACCESSed by another user and the data space was built at that time.

The user and server virtual machines are run in XC mode.

User Results

Data in a DIRCONTROL directory in a VM data space:

COMMAND	INSTR	N/S PGS	SHR PGS	TOT PGS	SPC OPS
UACCMCDS	19038	27	89	116	126
UMCDSXED	117297	61	126	187	263

Data in a FILECONTROL directory:

COMMAND	INSTR	N/S PGS	SHR PGS	TOT PGS	SPC OPS
UACC500	104171	41	93	134	176
UXED500	123184	61	127	188	248

Server Results

Data in a DIRCONTROL directory in a VM data space:

COMMAND	INSTR	N/S PGS	SHR PGS	TOT PGS	SPC OPS
SACCMCDS	17102	5	84	89	33
SMCDSXED	3606	5	48	53	36

Data in a FILECONTROL directory:

COMMAND	INSTR	N/S PGS	SHR PGS	TOT PGS	SPC OPS
SACC500	530604	7	191	198	401
SXED500	68890	5	139	144	187

Appendix B. SFS Counter Data

The SFS counts and timings in this appendix are provided to supplement the information provided for 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.

The QUERY FILEPOOL STATUS information was obtained for each SFS file pool server and the CRR recovery server. The counts and timings for each server were added together.

The first section in each table consists of the counters normalized by the number of commands (as determined by TPNS). The beginning values were subtracted from the ending values and divided by the number of commands in the interval. For each table, counts and timings which have a value of zero for all measurements shown in that table are not shown. A zero entry indicates that at least one occurrence was counted but the result of normalizing per command is so small that it rounds to zero. A description of the SFS counts and timings can be found in *VM/ESA 1.1 CMS Administration Reference*.

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 the glossary for definitions of these derived values.

PAGE REFERENCES, at the top of each column of counter values, shows the page number where the remaining measurement data for that run appears in the main body of this document. Multiple page numbers are shown in cases where that run is used in multiple places.

9021-720: CMS Regression / Software Modes / IPOLL

PAGE REFERENCES	121	60,96,121	121,167	60,96,121	121,163	207
RELEASE	ESA 1.0	ESA 1.0	ESA 1.1	ESA 1.1	ESA 1.1	ESA 1.1
RUN ID	Y63F480B	Y63F4809	Y64F480M	Y64F480X	Y64F480L	Y64F480V
PROCESSOR	9021-720	9021-720	9021-720	9021-720	9021-720	9021-720
REAL STORAGE	512M	512M	512M	512M	512M	512M
EXP. STORAGE	2048M	2048M	2048M	2048M	2048M	2048M
WORKLOAD	FS7B35R	FS7B35R	FS7B35R	FS7B35R	FS7B35R	FS7B35R
USERS	4800	4800	4800	4800	4800	4800
NORMALIZED BY COMMAND						
Close File Requests	0.3972	0.3967	0.3938	0.3938	0.3955	0.3929
Connect Requests	0.0032	0.0033	0.0046	0.0044	0.0044	0.0044
Delete File Requests	0.0953	0.0960	0.0959	0.0949	0.0950	0.0953
Lock Requests	0.0246	0.0248	0.0244	0.0246	0.0246	0.0245
Open File New Requests	0.0016	0.0016	0.0015	0.0015	0.0016	0.0016
Open File Read Requests	0.2273	0.2271	0.2251	0.2262	0.2279	0.2244
Open File Replace Requests	0.1454	0.1450	0.1462	0.1449	0.1447	0.1456
Open File Write Requests	0.0228	0.0231	0.0210	0.0212	0.0214	0.0213
Query File Pool Requests	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Query User Space Requests	0.0213	0.0215	0.0213	0.0214	0.0213	0.0214
Read File Requests	0.2282	0.2300	0.2314	0.2320	0.2306	0.2302
Refresh Directory Requests	0.0097	0.0099	0.0099	0.0096	0.0097	0.0095
Rename Requests	0.0050	0.0050	0.0049	0.0050	0.0050	0.0049
Unlock Requests	0.0247	0.0246	0.0245	0.0246	0.0244	0.0248
Write File Requests	0.1280	0.1288	0.1354	0.1342	0.1332	0.1347
Total File Pool Requests	1.3342	1.3373	1.3399	1.3383	1.3391	1.3356
File Pool Request Service Time	130.5655	127.0508	87.3487	87.5938	87.4313	88.0324
Local File Pool Requests	1.3342	1.3373	1.3399	1.3383	1.3391	1.3356
Begin LUWs	0.5020	0.5029	0.4981	0.4986	0.4980	0.4965
Agent Holding Time (msec)	168.1509	165.5363	119.4016	120.0613	121.4617	121.3625
SAC Calls	6.2634	6.2506	6.2558	6.2368	6.1914	6.2329
Catalog Lock Conflicts	0.0056	0.0060	0.0036	0.0032	0.0036	0.0029
Total Lock Conflicts	0.0056	0.0060	0.0036	0.0032	0.0036	0.0029
Lock Wait Time (msec)	0.3570	0.3489	0.1823	0.1459	0.1829	0.1718
File Blocks Read	0.9452	0.9509	0.9450	0.9475	0.9467	0.9426
File Blocks Written	0.5859	0.5864	0.6006	0.5953	0.5926	0.5978
Catalog Blocks Read	0.4854	0.4759	0.4735	0.4777	0.4654	0.4720
Catalog Blocks Written	0.2787	0.2750	0.2515	0.2518	0.2418	0.2500
Control Minidisk Blocks Read	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Control Minidisk Blocks Written	0.1039	0.1040	0.0654	0.0647	0.0631	0.0650
Log Blocks Written	0.4503	0.4569	0.4820	0.4840	0.4785	0.4833
Total DASD Block Transfers	2.8494	2.8492	2.8181	2.8209	2.7882	2.8107
BIO Requests to Read File Block	0.4966	0.4995	0.4937	0.4962	0.4968	0.4937
BIO Requests to Write File Blocks	0.2847	0.2853	0.2916	0.2900	0.2889	0.2922
BIO Requests to Read Catalog Blocks	0.4854	0.4759	0.4735	0.4777	0.4654	0.4720
BIO Requests to Write Catalog Blocks	0.2787	0.2750	0.2515	0.2518	0.2418	0.2500
BIO Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
BIO Requests to Write Cntrl Mdsk Blks	0.0038	0.0038	0.0022	0.0022	0.0022	0.0022
BIO Requests to Write Log Blocks	0.4503	0.4569	0.4820	0.4840	0.4785	0.4833
Total BIO Requests	1.9996	1.9965	1.9947	2.0019	1.9738	1.9934
Total BIO Request Time (msec)	42.3635	40.5999	36.9903	36.3979	36.4517	36.9693
I/O Requests to Read File Blocks	0.4788	0.4810	0.4536	0.4563	0.4576	0.4535
I/O Requests to Write File Blocks	0.3025	0.3060	0.3090	0.3080	0.3070	0.3095
I/O Requests to Read Catalog Blocks	0.4854	0.4759	0.4735	0.4777	0.4654	0.4720
I/O Requests to Write Catalog Blocks	0.2787	0.2750	0.2515	0.2518	0.2418	0.2500
I/O Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
I/O Requests to Write Cntrl Mdsk Blks	0.0072	0.0072	0.0041	0.0040	0.0040	0.0040
I/O Requests to Write Log Blocks	0.4503	0.4569	0.4820	0.4840	0.4785	0.4833
Total I/O Requests	2.0030	2.0021	1.9738	1.9818	1.9544	1.9724

PAGE REFERENCES	121	60,96,121	121,167	60,96,121	121,163	207
RELEASE	**ESA 1.0**	**ESA 1.0**	**ESA 1.1**	**ESA 1.1**	**ESA 1.1**	**ESA 1.1**
RUN ID	**Y63F480B**	**Y63F4809**	**Y64F480M**	**Y64F480X**	**Y64F480L**	**Y64F480V**
PROCESSOR	**9021-720**	**9021-720**	**9021-720**	**9021-720**	**9021-720**	**9021-720**
REAL STORAGE	**512M**	**512M**	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**	**2048M**	**2048M**
WORKLOAD	**FS7B35R**	**FS7B35R**	**FS7B35R**	**FS7B35R**	**FS7B35R**	**FS7B35R**
USERS	**4800**	**4800**	**4800**	**4800**	**4800**	**4800**
Get Logname Requests	0.0032	0.0033	0.0033	0.0032	0.0032	0.0032
Get LUWID Requests	0.0000	0.0000	0.0033	0.0032	0.0032	0.0032
Total CRR Requests	0.0032	0.0033	0.0066	0.0064	0.0064	0.0064
CRR Request Service Time (msec)	0.0017	0.0017	0.0602	0.0576	0.0579	0.0577
Log I/O Requests	0.0000	0.0000	0.0066	0.0064	0.0064	0.0063
DERIVED RESULTS						
Agents Held	28.4	27.9	20.3	20.4	20.6	20.5
Agents In-call	22.0	21.4	14.8	14.9	14.8	14.9
Avg LUW Time (msec)	335.0	329.2	239.7	240.8	243.9	244.4
Avg File Pool Request Time (msec)	97.9	95.0	65.2	65.5	65.3	65.9
Avg Lock Wait Time (msec)	63.8	58.2	50.6	45.6	50.8	59.2
SAC Calls / FP Request	4.69	4.67	4.67	4.66	4.62	4.67
Deadlocks (delta)	0	0	0	0	0	0
Rollbacks Due to Deadlock (delta)	0	1	0	0	0	0
Rollback Requests (delta)	0	0	0	0	0	0
LUW Rollbacks (delta)	0	1	0	0	0	0
Checkpoints Taken (delta)	202	198	126	114	122	115
Checkpoint Duration (sec)	5.2	5.1	5.1	5.2	5.2	5.1
Seconds Between Checkpoints	9.4	9.3	15.7	16.1	16.4	16.1
Checkpoint Utilization	55.3	54.7	32.5	32.6	31.6	32.1
BIO Request Time (msec)	21.19	20.34	18.54	18.18	18.47	18.55
Blocking Factor (Blocks/BIO)	1.42	1.43	1.41	1.41	1.41	1.41
Chaining Factor (Blocks/IO)	1.42	1.42	1.43	1.42	1.43	1.43

9021-720: Storage Constrained

PAGE REFERENCES	158	158	158	158	158	158
RELEASE	ESA 1.1	ESA 1.1	ESA 1.1	ESA 1.1	ESA 1.1	ESA 1.1
RUN ID	Y64F4808	Y64F480A	Y64F480D	Y64F480E	Y64F4809	Y64F480F
PROCESSOR	9021-720	9021-720	9021-720	9021-720	9021-720	9021-720
REAL/EXP. STORAGE	256M	256M	320M	320M	384M	512M
WORKLOAD	512M	768M	896M	1024M	1024M	2048M
USERS	FS7B35R	FS7B35R	FS7B35R	FS7B35R	FS7B35R	FS7B35R
	4800	4800	4800	4800	4800	4800
NORMALIZED BY COMMAND						
Close File Requests	0.3860	0.3980	0.4056	0.3941	0.3974	0.3971
Connect Requests	0.0041	0.0043	0.0045	0.0043	0.0043	0.0044
Delete File Requests	0.1033	0.0967	0.0970	0.0949	0.0955	0.0963
Lock Requests	0.0278	0.0247	0.0257	0.0248	0.0247	0.0247
Open File New Requests	0.0021	0.0016	0.0017	0.0016	0.0016	0.0016
Open File Read Requests	0.2026	0.2293	0.2352	0.2269	0.2288	0.2285
Open File Replace Requests	0.1608	0.1454	0.1463	0.1440	0.1455	0.1460
Open File Write Requests	0.0214	0.0213	0.0223	0.0215	0.0215	0.0210
Query File Pool Requests	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000
Query User Space Requests	0.0245	0.0207	0.0216	0.0212	0.0213	0.0215
Read File Requests	0.2402	0.2380	0.2414	0.2312	0.2321	0.2325
Refresh Directory Requests	0.0078	0.0091	0.0097	0.0095	0.0095	0.0095
Rename Requests	0.0051	0.0050	0.0050	0.0050	0.0049	0.0050
Unlock Requests	0.0279	0.0247	0.0257	0.0247	0.0245	0.0246
Write File Requests	0.1453	0.1366	0.1378	0.1329	0.1329	0.1348
Total File Pool Requests	1.3589	1.3553	1.3795	1.3365	1.3446	1.3475
File Pool Request Service Time	1693.8389	406.7864	391.6673	138.3310	129.3098	84.8561
Local File Pool Requests	1.3589	1.3553	1.3795	1.3365	1.3446	1.3475
Begin LUWs	0.4831	0.4959	0.5096	0.4972	0.5006	0.5025
Agent Holding Time (msec)	2925.5490	550.0576	506.4117	180.5697	170.2002	118.2676
SAC Calls	6.2168	6.1685	6.2886	6.1878	6.2533	6.2888
Catalog Lock Conflicts	0.0462	0.0164	0.0193	0.0048	0.0050	0.0032
Total Lock Conflicts	0.0462	0.0164	0.0193	0.0048	0.0050	0.0032
Lock Wait Time (msec)	164.9504	34.0655	24.2613	0.5111	0.3433	0.1731
File Blocks Read	0.9617	0.9694	0.9868	0.9471	0.9498	0.9506
File Blocks Written	0.6461	0.6007	0.6064	0.5904	0.5940	0.5992
Catalog Blocks Read	0.8181	0.6806	0.6377	0.5004	0.4990	0.4917
Catalog Blocks Written	0.3785	0.3242	0.3089	0.2540	0.2567	0.2528
Control Minidisk Blocks Read	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Control Minidisk Blocks Written	0.0737	0.0643	0.0654	0.0636	0.0640	0.0642
Log Blocks Written	0.3984	0.4203	0.4076	0.4467	0.4536	0.4862
Total DASD Block Transfers	3.2765	3.0596	3.0128	2.8021	2.8171	2.8447
BIO Requests to Read File Block	0.4897	0.5113	0.5224	0.5003	0.5011	0.4985
BIO Requests to Write File Blocks	0.3200	0.2945	0.2988	0.2906	0.2900	0.2903
BIO Requests to Read Catalog Blocks	0.8181	0.6806	0.6377	0.5004	0.4990	0.4917
BIO Requests to Write Catalog Blocks	0.3785	0.3242	0.3089	0.2540	0.2567	0.2528
BIO Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
BIO Requests to Write Cntrl Mdsk Blks	0.0024	0.0022	0.0022	0.0022	0.0022	0.0022
BIO Requests to Write Log Blocks	0.3984	0.4203	0.4076	0.4467	0.4536	0.4862
Total BIO Requests	2.4072	2.2332	2.1777	1.9942	2.0027	2.0218
Total BIO Request Time (msec)	707.2909	174.1410	173.5636	70.6825	67.8573	36.7579
I/O Requests to Read File Blocks	0.4798	0.5106	0.5230	0.4980	0.4964	0.4602
I/O Requests to Write File Blocks	0.3414	0.3161	0.3211	0.3106	0.3080	0.3076
I/O Requests to Read Catalog Blocks	0.8181	0.6806	0.6377	0.5004	0.4990	0.4917
I/O Requests to Write Catalog Blocks	0.3785	0.3242	0.3089	0.2540	0.2567	0.2528
I/O Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
I/O Requests to Write Cntrl Mdsk Blks	0.0045	0.0040	0.0041	0.0040	0.0040	0.0041
I/O Requests to Write Log Blocks	0.3984	0.4203	0.4076	0.4467	0.4536	0.4862
Total I/O Requests	2.4206	2.2559	2.2025	2.0137	2.0178	2.0026

PAGE REFERENCES RELEASE RUN ID PROCESSOR REAL/EXP. STORAGE WORKLOAD USERS	158 ESA 1.1 Y64F4808 9021-720 256M 512M FS7B35R 4800	158 ESA 1.1 Y64F480A 9021-720 256M 768M FS7B35R 4800	158 ESA 1.1 Y64F480D 9021-720 320M 896M FS7B35R 4800	158 ESA 1.1 Y64F480E 9021-720 320M 1024M FS7B35R 4800	158 ESA 1.1 Y64F4809 9021-720 384M 1024M FS7B35R 4800	158 ESA 1.1 Y64F480F 9021-720 512M 2048M FS7B35R 4800
Get Logname Requests	0.0026	0.0030	0.0032	0.0032	0.0031	0.0032
Get LUWID Requests	0.0026	0.0030	0.0032	0.0032	0.0031	0.0032
Total CRR Requests	0.0051	0.0060	0.0065	0.0063	0.0063	0.0063
CRR Request Service Time (msec)	1.8546	0.2701	0.0817	0.0654	0.0613	0.0575
Log I/O Requests	0.0034	0.0058	0.0064	0.0063	0.0063	0.0063
DERIVED RESULTS						
Agents Held	152.6	72.6	71.7	30.5	28.8	20.1
Agents In-call	88.3	53.7	55.4	23.3	21.9	14.4
Avg LUW Time (msec)	6055.8	1109.2	993.7	363.2	340.0	235.4
Avg File Pool Request Time (msec)	1246.5	300.1	283.9	103.5	96.2	63.0
Avg Lock Wait Time (msec)	3570.4	2077.2	1257.1	106.5	68.7	54.1
SAC Calls / FP Request	4.57	4.55	4.56	4.63	4.65	4.67
Deadlocks (delta)	0	3	0	0	0	0
Rollbacks Due to Deadlock (delta)	0	0	0	0	0	0
Rollback Requests (delta)	0	0	0	0	0	0
LUW Rollbacks (delta)	12	14	8	0	1	0
Checkpoints Taken (delta)	35	85	94	110	112	117
Checkpoint Duration (sec)	16.6	6.4	6.6	5.3	5.1	4.9
Seconds Between Checkpoints	47.4	20.8	19.1	16.4	16.1	15.9
Checkpoint Utilization	35.2	30.9	34.3	32.2	32.0	31.0
BIO Request Time (msec)	293.82	77.98	79.70	35.44	33.88	18.18
Blocking Factor (Blocks/BIO)	1.36	1.37	1.38	1.41	1.41	1.41
Chaining Factor (Blocks/IO)	1.35	1.36	1.37	1.39	1.40	1.42

9021-720: Data Spaces

PAGE REFERENCES	163	163	163	167
RELEASE	**ESA 1.1**	**ESA 1.1**	**ESA 1.1**	**ESA 1.1**
RUN ID	**Y64M480J**	**Y64M480K**	**Y64M480O**	**Y64M480P**
PROCESSOR	**9021-720**	**9021-720**	**9021-720**	**9021-720**
REAL STORAGE	**512M**	**512M**	**512M**	**512M**
EXP. STORAGE	**2048M**	**2048M**	**2048M**	**2048M**
WORKLOAD	**FS7B100R**	**FS7B35R**	**FS7B100R**	**FS7B100R**
USERS	**4800**	**4800**	**4800**	**4800**
NORMALIZED BY COMMAND				
Close File Requests	0.3973	0.3984	0.7154	0.3949
Commit Requests	0.0344	0.0000	0.0690	0.0346
Connect Requests	0.0076	0.0046	0.0074	0.0076
Delete File Requests	0.0963	0.0963	0.0949	0.0957
Lock Requests	0.0249	0.0247	0.0247	0.0247
Open File New Requests	0.0016	0.0016	0.0015	0.0015
Open File Read Requests	0.2280	0.2292	0.5483	0.2269
Open File Replace Requests	0.1462	0.1464	0.1438	0.1453
Open File Write Requests	0.0215	0.0214	0.0217	0.0212
Query File Pool Requests	0.0000	0.0000	0.0000	0.0000
Query User Space Requests	0.0218	0.0215	0.0212	0.0216
Read File Requests	0.2344	0.2329	0.3533	0.2326
Refresh Directory Requests	0.0225	0.0098	0.0218	0.0221
Rename Requests	0.0049	0.0050	0.0050	0.0050
Rollback Requests	0.0000	0.0000	0.0000	0.0000
Unlock Requests	0.0246	0.0246	0.0246	0.0248
Write File Requests	0.1349	0.1353	0.1331	0.1351
Total File Pool Requests	1.4008	1.3516	2.1859	1.3935
File Pool Request Service Time	83.1895	85.5118	96.9508	84.4537
Local File Pool Requests	1.4008	1.3516	2.1859	1.3935
Begin LUWs	0.5217	0.5044	0.8402	0.5164
Agent Holding Time (msec)	117.1895	118.8361	153.0763	116.3201
SAC Calls	6.4249	6.3116	9.2869	6.3194
Catalog Lock Conflicts	0.0031	0.0032	0.0039	0.0033
Total Lock Conflicts	0.0031	0.0032	0.0039	0.0033
Lock Wait Time (msec)	0.1874	0.1765	0.2218	0.1648
File Blocks Read	0.9570	0.9533	1.5868	0.9495
File Blocks Written	0.6006	0.6008	0.5913	0.5987
Catalog Blocks Read	0.4776	0.4763	0.4824	0.4590
Catalog Blocks Written	0.2529	0.2543	0.2538	0.2401
Control Minidisk Blocks Read	0.0000	0.0000	0.0000	0.0000
Control Minidisk Blocks Written	0.0644	0.0657	0.0633	0.0637
Log Blocks Written	0.4896	0.4879	0.4818	0.4817
Total DASD Block Transfers	2.8421	2.8383	3.4594	2.7926
BIO Requests to Read File Block	0.5004	0.4997	0.9462	0.4968
BIO Requests to Write File Blocks	0.2911	0.2914	0.2890	0.2895
BIO Requests to Read Catalog Blocks	0.4776	0.4763	0.4824	0.4590
BIO Requests to Write Catalog Blocks	0.2529	0.2543	0.2538	0.2401
BIO Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000	0.0000
BIO Requests to Write Cntrl Mdsk Blks	0.0022	0.0023	0.0022	0.0022
BIO Requests to Write Log Blocks	0.4896	0.4879	0.4818	0.4817
Total BIO Requests	2.0138	2.0118	2.4554	1.9694
Total BIO Request Time (msec)	34.6015	34.4338	39.0680	35.6757

PAGE REFERENCES RELEASE RUN ID PROCESSOR REAL STORAGE EXP. STORAGE WORKLOAD USERS	163 ESA 1.1 Y64M480J 9021-720 512M 2048M FS7B100R 4800	163 ESA 1.1 Y64M480K 9021-720 512M 2048M FS7B35R 4800	163 ESA 1.1 Y64M480O 9021-720 512M 2048M FS7B100R 4800	167 ESA 1.1 Y64M480P 9021-720 512M 2048M FS7B100R 4800
I/O Requests to Read File Blocks	0.4574	0.4547	0.8737	0.4557
I/O Requests to Write File Blocks	0.3085	0.3088	0.3067	0.3076
I/O Requests to Read Catalog Blocks	0.4776	0.4763	0.4824	0.4590
I/O Requests to Write Catalog Blocks	0.2529	0.2543	0.2538	0.2401
I/O Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000	0.0000
I/O Requests to Write Cntrl Mdsk Blks	0.0041	0.0041	0.0039	0.0040
I/O Requests to Write Log Blocks	0.4896	0.4879	0.4818	0.4817
Total I/O Requests	1.9901	1.9862	2.4024	1.9480
Get Logname Requests	0.0032	0.0033	0.0031	0.0031
Get LUWID Requests	0.0032	0.0033	0.0031	0.0031
Total CRR Requests	0.0064	0.0066	0.0062	0.0063
CRR Request Service Time (msec)	0.0575	0.0586	0.0603	0.0565
Log I/O Requests	0.0064	0.0065	0.0062	0.0063
DERIVED RESULTS				
Agents Held	19.9	20.1	25.9	19.7
Agents In-call	14.1	14.5	16.4	14.3
Avg LUW Time (msec)	224.6	235.6	182.2	225.3
Avg File Pool Request Time (msec)	59.4	63.3	44.4	60.6
Avg Lock Wait Time (msec)	60.5	55.2	56.9	49.9
SAC Calls / FP Request	4.59	4.67	4.25	4.53
Deadlocks (delta)	0	0	0	0
Rollbacks Due to Deadlock (delta)	0	0	0	0
Rollback Requests (delta)	0	0	0	0
LUW Rollbacks (delta)	0	0	0	0
Checkpoints Taken (delta)	113	115	110	112
Checkpoint Duration (sec)	5.0	5.1	5.3	5.0
Seconds Between Checkpoints	15.9	15.7	16.3	16.1
Checkpoint Utilization	31.5	32.3	32.5	31.0
BIO Request Time (msec)	17.18	17.12	15.91	18.12
Blocking Factor (Blocks/BIO)	1.41	1.41	1.41	1.42
Chaining Factor (Blocks/IO)	1.43	1.43	1.44	1.43

9021-580: CMS Regression / VM Storage Considerations

PAGE REFERENCES RELEASE RUN ID PROCESSOR REAL STORAGE EXP. STORAGE WORKLOAD USERS	64 ESA 1.0 Y33F2642 9021-580 256M 1024M FS7B35R 2640	64,113,117 ESA 1.1 Y34F2644 9021-580 256M 1024M FS7B35R 2640	113 ESA 1.1 Y34F2645 9021-580 256M 1024M FS7B35R 2640	117 ESA 1.1 Y34F2647 9021-580 256M 1024M FS7B35R 2640	117 ESA 1.1 Y34F2648 9021-580 256M 1024M FS7B35R 2640
NORMALIZED BY COMMAND					
Close File Requests	0.4010	0.3983	0.3953	0.3968	0.3922
Connect Requests	0.0032	0.0040	0.0042	0.0041	0.0042
Delete File Requests	0.0931	0.0930	0.0938	0.0935	0.0927
Lock Requests	0.0248	0.0247	0.0245	0.0248	0.0247
Open File New Requests	0.0015	0.0015	0.0014	0.0015	0.0014
Open File Read Requests	0.2356	0.2347	0.2308	0.2329	0.2297
Open File Replace Requests	0.1407	0.1405	0.1414	0.1412	0.1396
Open File Write Requests	0.0235	0.0213	0.0216	0.0215	0.0214
Query File Pool Requests	0.0000	0.0000	0.0000	0.0000	0.0000
Query User Space Requests	0.0201	0.0203	0.0206	0.0204	0.0200
Read File Requests	0.2172	0.2148	0.2144	0.2179	0.2165
Refresh Directory Requests	0.0096	0.0091	0.0097	0.0093	0.0096
Rename Requests	0.0049	0.0048	0.0048	0.0049	0.0048
Unlock Requests	0.0246	0.0245	0.0243	0.0245	0.0245
Write File Requests	0.1208	0.1249	0.1253	0.1274	0.1267
Total File Pool Requests	1.3204	1.3165	1.3121	1.3206	1.3080
File Pool Request Service Time	219.1791	94.8351	98.9756	97.6708	96.9892
Local File Pool Requests	1.3204	1.3165	1.3121	1.3206	1.3080
Begin LUWs	0.5050	0.4982	0.4998	0.5016	0.4971
Agent Holding Time (msec)	325.1333	158.0850	162.5474	163.4125	162.5739
SAC Calls	6.2125	6.0973	6.1762	6.1999	6.1359
Catalog Lock Conflicts	0.0117	0.0047	0.0056	0.0057	0.0051
Total Lock Conflicts	0.0117	0.0047	0.0056	0.0057	0.0051
Lock Wait Time (msec)	1.6047	0.4583	0.4114	0.5419	0.3562
File Blocks Read	0.9374	0.9240	0.9186	0.9312	0.9239
File Blocks Written	0.5606	0.5681	0.5700	0.5737	0.5696
Catalog Blocks Read	0.4097	0.3868	0.4017	0.3987	0.3980
Catalog Blocks Written	0.2479	0.2021	0.2129	0.2147	0.2114
Control Minidisk Blocks Written	0.1142	0.0691	0.0699	0.0721	0.0703
Log Blocks Written	0.3956	0.4651	0.4706	0.4732	0.4677
Total DASD Block Transfers	2.6655	2.6152	2.6437	2.6635	2.6408
BIO Requests to Read File Block	0.4947	0.4847	0.4814	0.4876	0.4832
BIO Requests to Write File Blocks	0.2724	0.2763	0.2778	0.2796	0.2779
BIO Requests to Read Catalog Blocks	0.4097	0.3868	0.4017	0.3987	0.3980
BIO Requests to Write Catalog Blocks	0.2479	0.2021	0.2129	0.2147	0.2114
BIO Requests to Write Cntrl Mdsk Blks	0.0037	0.0021	0.0021	0.0022	0.0021
BIO Requests to Write Log Blocks	0.3956	0.4651	0.4706	0.4732	0.4677
Total BIO Requests	1.8240	1.8172	1.8465	1.8560	1.8403
Total BIO Request Time (msec)	52.4402	33.3351	34.1284	32.6152	34.0485

PAGE REFERENCES RELEASE RUN ID PROCESSOR REAL STORAGE EXP. STORAGE WORKLOAD USERS	64 ESA 1.0 Y33F2642 9021-580 256M 1024M FS7B35R 2640	64,113,117 ESA 1.1 Y34F2644 9021-580 256M 1024M FS7B35R 2640	113 ESA 1.1 Y34F2645 9021-580 256M 1024M FS7B35R 2640	117 ESA 1.1 Y34F2647 9021-580 256M 1024M FS7B35R 2640	117 ESA 1.1 Y34F2648 9021-580 256M 1024M FS7B35R 2640
I/O Requests to Read File Blocks	0.4717	0.4520	0.4479	0.4493	0.4485
I/O Requests to Write File Blocks	0.2882	0.2925	0.2943	0.2955	0.2944
I/O Requests to Read Catalog Blocks	0.4097	0.3868	0.4017	0.3987	0.3980
I/O Requests to Write Catalog Blocks	0.2479	0.2021	0.2129	0.2147	0.2114
I/O Requests to Write Cntrl Mdsk Blks	0.0076	0.0041	0.0043	0.0044	0.0043
I/O Requests to Write Log Blocks	0.3956	0.4651	0.4706	0.4732	0.4677
Total I/O Requests	1.8206	1.8027	1.8316	1.8358	1.8243
Get Logname Requests	0.0032	0.0030	0.0032	0.0031	0.0032
Get LUWID Requests	0.0000	0.0030	0.0032	0.0031	0.0032
Total CRR Requests	0.0032	0.0061	0.0065	0.0062	0.0064
CRR Request Service Time (msec)	0.0015	0.0467	0.0479	0.0461	0.0495
Log I/O Requests	0.0000	0.0061	0.0065	0.0062	0.0064
DERIVED RESULTS					
Agents Held	30.8	15.1	15.6	15.6	15.5
Agents In-call	20.8	9.0	9.5	9.3	9.2
Avg LUW Time (msec)	643.8	317.3	325.2	325.8	327.0
Avg File Pool Request Time (msec)	166.0	72.0	75.4	74.0	74.2
Avg Lock Wait Time (msec)	137.2	97.5	73.5	95.1	69.8
SAC Calls / FP Request	4.71	4.63	4.71	4.69	4.69
Deadlocks (delta)	0	0	0	0	0
Rollbacks Due to Deadlock (delta)	0	0	0	0	0
Rollback Requests (delta)	0	0	0	0	0
LUW Rollbacks (delta)	0	0	0	0	0
Checkpoints Taken (delta)	106	60	61	63	61
Checkpoint Duration (sec)	6.6	5.4	5.6	5.6	5.5
Seconds Between Checkpoints	17.0	30.0	29.5	28.6	29.5
Checkpoint Utilization	38.7	18.0	18.9	19.8	18.6
BIO Request Time (msec)	28.75	18.34	18.48	17.57	18.50
Blocking Factor (Blocks/BIO)	1.46	1.44	1.43	1.44	1.43
Chaining Factor (Blocks/IO)	1.46	1.45	1.44	1.45	1.45

9121-480: Tuning / CMS Regression

PAGE REFERENCES RELEASE RUN ID PROCESSOR REAL STORAGE EXP. STORAGE WORKLOAD USERS	233 ESA 1.0 L23F1480 9121-480 192M 64M FS7B35R 1480	230 ESA 1.0 L23F1481 9121-480 192M 64M FS7B35R 1480	230 ESA 1.0 L23F1482 9121-480 256M 0M FS7B35R 1480	72,233 ESA 1.0 L23F1484 9121-480 192M 64M FS7B35R 1480	72 ESA 1.1 L24F1480 9121-480 192M 64M FS7B35R 1480
NORMALIZED BY COMMAND					
Close File Requests	0.3549	0.3953	0.3934	0.3890	0.3910
Connect Requests	0.0028	0.0031	0.0033	0.0033	0.0044
Delete File Requests	0.0909	0.0965	0.0959	0.0960	0.0942
Lock Requests	0.0226	0.0243	0.0246	0.0248	0.0246
Open File New Requests	0.0014	0.0019	0.0018	0.0018	0.0019
Open File Read Requests	0.1941	0.2240	0.2234	0.2170	0.2228
Open File Replace Requests	0.1365	0.1466	0.1456	0.1476	0.1443
Open File Write Requests	0.0204	0.0227	0.0229	0.0229	0.0220
Query File Pool Requests	0.0001	0.0000	0.0000	0.0000	0.0000
Query User Space Requests	0.0207	0.0218	0.0220	0.0220	0.0213
Read File Requests	0.2210	0.2356	0.2349	0.2352	0.2339
Refresh Directory Requests	0.0090	0.0092	0.0098	0.0098	0.0097
Rename Requests	0.0047	0.0041	0.0042	0.0041	0.0041
Unlock Requests	0.0218	0.0244	0.0243	0.0247	0.0249
Write File Requests	0.1178	0.1286	0.1271	0.1292	0.1344
Total File Pool Requests	1.2187	1.3381	1.3332	1.3277	1.3334
File Pool Request Service Time	4618.0236	212.6156	210.1902	224.2788	151.6628
Local File Pool Requests	1.2187	1.3381	1.3332	1.3277	1.3334
Begin LUWs	0.4379	0.4881	0.4886	0.4820	0.4884
Agent Holding Time (msec)	na	362.2621	382.9755	395.3427	249.2847
SAC Calls	5.5483	6.0859	6.0662	6.0536	6.1219
Catalog Lock Conflicts	0.0546	0.0075	0.0068	0.0079	0.0059
Total Lock Conflicts	0.0546	0.0075	0.0068	0.0079	0.0059
Lock Wait Time (msec)	187.1798	2.4945	1.8754	2.9857	1.7871
File Blocks Read	0.8853	0.9659	0.9633	0.9588	0.9544
File Blocks Written	0.5442	0.5906	0.5829	0.5936	0.5970
Catalog Blocks Read	1.0286	0.3416	0.3451	0.3524	0.3524
Catalog Blocks Written	0.4709	0.1957	0.1942	0.1967	0.1749
Control Minidisk Blocks Read	0.0000	0.0000	0.0000	0.0000	0.0000
Control Minidisk Blocks Written	0.0851	0.0909	0.0892	0.0898	0.0524
Log Blocks Written	0.1146	0.4798	0.4903	0.4788	0.5169
Total DASD Block Transfers	3.1287	2.6644	2.6650	2.6700	2.6480
BIO Requests to Read File Block	0.4657	0.5036	0.5053	0.4961	0.4941
BIO Requests to Write File Blocks	0.2684	0.2895	0.2871	0.2903	0.2924
BIO Requests to Read Catalog Blocks	1.0286	0.3416	0.3451	0.3524	0.3524
BIO Requests to Write Catalog Blocks	0.4709	0.1957	0.1942	0.1967	0.1749
BIO Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000	0.0000	0.0000
BIO Requests to Write Cntrl Mdsk Blks	0.0034	0.0038	0.0037	0.0037	0.0021
BIO Requests to Write Log Blocks	0.1146	0.4798	0.4903	0.4788	0.5169
Total BIO Requests	2.3517	1.8140	1.8257	1.8180	1.8327
Total BIO Request Time (msec)	2273.1859	68.3485	74.5859	70.5469	54.9035

PAGE REFERENCES	233	230	230	72,233	72
RELEASE	**ESA 1.0**	**ESA 1.0**	**ESA 1.0**	**ESA 1.0**	**ESA 1.1**
RUN ID	**L23F1480**	**L23F1481**	**L23F1482**	**L23F1484**	**L24F1480**
PROCESSOR	**9121-480**	**9121-480**	**9121-480**	**9121-480**	**9121-480**
REAL STORAGE	**192M**	**192M**	**256M**	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**	**0M**	**64M**	**64M**
WORKLOAD	**FS7B35R**	**FS7B35R**	**FS7B35R**	**FS7B35R**	**FS7B35R**
USERS	**1480**	**1480**	**1480**	**1480**	**1480**
I/O Requests to Read File Blocks	0.4462	0.4902	0.5182	0.4837	0.4665
I/O Requests to Write File Blocks	0.2833	0.3039	0.2996	0.3053	0.3078
I/O Requests to Read Catalog Blocks	1.0286	0.3416	0.3451	0.3524	0.3524
I/O Requests to Write Catalog Blocks	0.4709	0.1957	0.1942	0.1967	0.1749
I/O Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000	0.0000	0.0000
I/O Requests to Write Cntrl Mdsk Blks	0.0074	0.0081	0.0080	0.0080	0.0045
I/O Requests to Write Log Blocks	0.1146	0.4798	0.4903	0.4788	0.5169
Total I/O Requests	2.3510	1.8192	1.8553	1.8249	1.8229
Get Logname Requests	0.0030	0.0031	0.0033	0.0033	0.0032
Get LUWID Requests	0.0000	0.0000	0.0000	0.0000	0.0032
Total CRR Requests	0.0030	0.0031	0.0033	0.0033	0.0064
CRR Request Service Time (msec)	11.1715	0.0017	0.0018	0.0019	0.1071
Log I/O Requests	0.0000	0.0000	0.0000	0.0000	0.0064
DERIVED RESULTS					
Agents Held	na	19.0	20.1	20.6	13.2
Agents In-call	88.6	11.2	11.0	11.7	8.0
Avg LUW Time (msec)	na	742.2	783.8	820.2	510.4
Avg File Pool Request Time (msec)	3789.3	158.9	157.7	168.9	113.7
Avg Lock Wait Time (msec)	3428.2	332.6	275.8	377.9	302.9
SAC Calls / FP Request	4.55	4.55	4.55	4.56	4.59
Deadlocks (delta)	2	0	0	0	1
Rollbacks Due to Deadlock (delta)	0	0	0	0	0
Rollback Requests (delta)	0	0	0	0	0
LUW Rollbacks (delta)	8	1	0	0	0
Checkpoints Taken (delta)	19	60	59	58	33
Checkpoint Duration (sec)	16.4	7.9	7.7	7.8	10.0
Seconds Between Checkpoints	91.7	30.3	31.0	31.1	54.3
Checkpoint Utilization	18.0	26.0	24.9	25.1	18.3
BIO Request Time (msec)	966.61	37.68	40.85	38.80	29.96
Blocking Factor (Blocks/BIO)	1.33	1.47	1.46	1.47	1.44
Chaining Factor (Blocks/IO)	1.33	1.46	1.44	1.46	1.45

9121-480: Data Spaces

PAGE REFERENCES	171	171	171
RELEASE	**ESA 1.1**	**ESA 1.1**	**ESA 1.1**
RUN ID	**L24F1481**	**L24F1482**	**L24F1483**
PROCESSOR	**9121-480**	**9121-480**	**9121-480**
REAL STORAGE	**192M**	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**	**64M**
WORKLOAD	**FS7B100R**	**FS7B35R**	**FS7B35R**
USERS	**1480**	**1480**	**1480**
NORMALIZED BY COMMAND			
Close File Requests	0.3916	0.3903	0.3916
Commit Requests	0.0338	0.0000	0.0000
Connect Requests	0.0078	0.0046	0.0044
Delete File Requests	0.0968	0.0949	0.0945
Lock Requests	0.0247	0.0246	0.0247
Open File New Requests	0.0018	0.0018	0.0018
Open File Read Requests	0.2215	0.2218	0.2252
Open File Replace Requests	0.1469	0.1454	0.1433
Open File Write Requests	0.0215	0.0210	0.0218
Query File Pool Requests	0.0000	0.0000	0.0000
Query User Space Requests	0.0218	0.0214	0.0216
Read File Requests	0.2347	0.2356	0.2375
Refresh Directory Requests	0.0228	0.0100	0.0094
Rename Requests	0.0042	0.0042	0.0041
Unlock Requests	0.0247	0.0246	0.0247
Write File Requests	0.1356	0.1327	0.1332
Total File Pool Requests	1.3904	1.3328	1.3380
File Pool Request Service Time	129.1871	167.0461	148.6991
Local File Pool Requests	1.3904	1.3328	1.3380
Begin LUWs	0.5066	0.4829	0.4919
Agent Holding Time (msec)	230.7265	268.4838	242.5336
SAC Calls	6.2119	6.0398	6.1484
Catalog Lock Conflicts	0.0045	0.0068	0.0054
Total Lock Conflicts	0.0045	0.0068	0.0054
Lock Wait Time (msec)	0.8461	2.0062	1.2115
File Blocks Read	0.9544	0.9568	0.9659
File Blocks Written	0.6028	0.5919	0.5914
Catalog Blocks Read	0.3465	0.3499	0.3519
Catalog Blocks Written	0.1703	0.1678	0.1847
Control Minidisk Blocks Read	0.0000	0.0000	0.0000
Control Minidisk Blocks Written	0.0539	0.0518	0.0558
Log Blocks Written	0.5171	0.5028	0.5204
Total DASD Block Transfers	2.6450	2.6210	2.6701
BIO Requests to Read File Block	0.4941	0.4946	0.5016
BIO Requests to Write File Blocks	0.2956	0.2912	0.2905
BIO Requests to Read Catalog Blocks	0.3465	0.3499	0.3519
BIO Requests to Write Catalog Blocks	0.1703	0.1678	0.1847
BIO Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000
BIO Requests to Write Cntrl Mdsk Blks	0.0021	0.0020	0.0022
BIO Requests to Write Log Blocks	0.5171	0.5028	0.5204
Total BIO Requests	1.8258	1.8083	1.8514
Total BIO Request Time (msec)	45.7297	56.6273	51.8401

PAGE REFERENCES	**171**	**171**	**171**
RELEASE	**ESA 1.1**	**ESA 1.1**	**ESA 1.1**
RUN ID	**L24F1481**	**L24F1482**	**L24F1483**
PROCESSOR	**9121-480**	**9121-480**	**9121-480**
REAL STORAGE	**192M**	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**	**64M**
WORKLOAD	**FS7B100R**	**FS7B35R**	**FS7B35R**
USERS	**1480**	**1480**	**1480**
I/O Requests to Read File Blocks	0.4661	0.4665	0.4749
I/O Requests to Write File Blocks	0.3099	0.3077	0.3055
I/O Requests to Read Catalog Blocks	0.3465	0.3499	0.3519
I/O Requests to Write Catalog Blocks	0.1703	0.1678	0.1847
I/O Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0000
I/O Requests to Write Cntrl Mdsk Blks	0.0046	0.0044	0.0048
I/O Requests to Write Log Blocks	0.5171	0.5028	0.5204
Total I/O Requests	1.8145	1.7992	1.8422
Get Logname Requests	0.0033	0.0033	0.0031
Get LUWID Requests	0.0033	0.0033	0.0031
Total CRR Requests	0.0065	0.0067	0.0063
CRR Request Service Time (msec)	0.1069	0.1179	0.1058
Log I/O Requests	0.0065	0.0067	0.0063
DERIVED RESULTS			
Agents Held	12.2	14.1	12.8
Agents In-call	6.8	8.7	7.8
Avg LUW Time (msec)	455.4	556.0	493.1
Avg File Pool Request Time (msec)	92.9	125.3	111.1
Avg Lock Wait Time (msec)	188.0	295.0	224.4
SAC Calls / FP Request	4.47	4.53	4.60
Deadlocks (delta)	0	0	0
Rollbacks Due to Deadlock (delta)	0	0	0
Rollback Requests (delta)	0	0	0
LUW Rollbacks (delta)	0	0	0
Checkpoints Taken (delta)	34	32	35
Checkpoint Duration (sec)	9.2	10.5	10.1
Seconds Between Checkpoints	52.9	56.2	51.5
Checkpoint Utilization	17.3	18.8	19.6
BIO Request Time (msec)	25.05	31.32	28.00
Blocking Factor (Blocks/BIO)	1.45	1.45	1.44
Chaining Factor (Blocks/IO)	1.46	1.46	1.45

9121-320: CMS Regression

PAGE REFERENCES	**80**	**80,105**
RELEASE	**ESA 1.0**	**ESA 1.1**
RUN ID	**L13F0771**	**L14F0770**
PROCESSOR	**9121-480**	**9121-480**
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
WORKLOAD	**FS7B35R**	**FS7B35R**
USERS	**770**	**770**
NORMALIZED BY COMMAND		
Close File Requests	0.3980	0.4000
Connect Requests	0.0031	0.0044
Delete File Requests	0.0944	0.0950
Lock Requests	0.0248	0.0245
Open File New Requests	0.0018	0.0019
Open File Read Requests	0.2300	0.2319
Open File Replace Requests	0.1440	0.1453
Open File Write Requests	0.0230	0.0215
Query File Pool Requests	0.0000	0.0000
Query User Space Requests	0.0215	0.0212
Read File Requests	0.2363	0.2387
Refresh Directory Requests	0.0093	0.0097
Rename Requests	0.0042	0.0041
Unlock Requests	0.0249	0.0246
Write File Requests	0.1277	0.1339
Total File Pool Requests	1.3431	1.3566
File Pool Request Service Time	148.8852	95.8809
Local File Pool Requests	1.3431	1.3566
Begin LUWs	0.5007	0.5029
Agent Holding Time (msec)	412.6111	240.2037
SAC Calls	6.2334	6.3070
Catalog Lock Conflicts	0.0049	0.0026
Total Lock Conflicts	0.0049	0.0026
Lock Wait Time (msec)	0.4000	0.1642
File Blocks Read	0.9737	0.9736
File Blocks Written	0.5830	0.5955
Catalog Blocks Read	0.3844	0.3695
Catalog Blocks Written	0.2196	0.1975
Control Minidisk Blocks Read	0.0000	0.0000
Control Minidisk Blocks Written	0.0936	0.0581
Log Blocks Written	0.5205	0.5423
Total DASD Block Transfers	2.7747	2.7365
BIO Requests to Read File Block	0.5150	0.5113
BIO Requests to Write File Blocks	0.2866	0.2933
BIO Requests to Read Catalog Blocks	0.3844	0.3695
BIO Requests to Write Catalog Blocks	0.2196	0.1975
BIO Requests to Read Cntrl Mdsk Blks	0.0000	0.0000
BIO Requests to Write Cntrl Mdsk Blks	0.0038	0.0023
BIO Requests to Write Log Blocks	0.5205	0.5423
Total BIO Requests	1.9298	1.9163
Total BIO Request Time (msec)	37.6304	31.2546

PAGE REFERENCES	**80**	**80,105**
RELEASE	**ESA 1.0**	**ESA 1.1**
RUN ID	**L13F0771**	**L14F0770**
PROCESSOR	**9121-480**	**9121-480**
REAL STORAGE	**192M**	**192M**
EXP. STORAGE	**64M**	**64M**
WORKLOAD	**FS7B35R**	**FS7B35R**
USERS	**770**	**770**
I/O Requests to Read File Blocks	0.4878	0.4631
I/O Requests to Write File Blocks	0.3029	0.3111
I/O Requests to Read Catalog Blocks	0.3844	0.3695
I/O Requests to Write Catalog Blocks	0.2196	0.1975
I/O Requests to Read Cntrl Mdsk Blks	0.0000	0.0000
I/O Requests to Write Cntrl Mdsk Blks	0.0082	0.0049
I/O Requests to Write Log Blocks	0.5205	0.5423
Total I/O Requests	1.9234	1.8884
Get Logname Requests	0.0031	0.0032
Get LUWID Requests	0.0000	0.0032
Total CRR Requests	0.0031	0.0064
CRR Request Service Time (msec)	0.0016	0.1092
Log I/O Requests	0.0000	0.0064
DERIVED RESULTS		
Agents Held	11.2	6.6
Agents In-call	4.0	2.6
Avg LUW Time (msec)	824.1	477.6
Avg File Pool Request Time (msec)	110.9	70.7
Avg Lock Wait Time (msec)	81.6	63.2
SAC Calls / FP Request	4.64	4.65
Deadlocks (delta)	0	0
Rollbacks Due to Deadlock (delta)	0	0
Rollback Requests (delta)	0	0
LUW Rollbacks (delta)	0	0
Checkpoints Taken (delta)	31	19
Checkpoint Duration (sec)	6.7	7.4
Seconds Between Checkpoints	58.1	95.2
Checkpoint Utilization	11.6	7.8
BIO Request Time (msec)	19.50	16.31
Blocking Factor (Blocks/BIO)	1.44	1.43
Chaining Factor (Blocks/IO)	1.44	1.45

9221-170: CMS Regression

PAGE REFERENCES	91,108	91,108,126	126,174	174
RELEASE	ESA 1.0	ESA 1.1	ESA 1.1	ESA 1.1
RUN ID	H17F0241	H14F0241	H14F0242	H14M0241
PROCESSOR	9221-170	9221-170	9221-170	9221-170
REAL STORAGE	64M	48M	48M	48M
EXP. STORAGE	0M	16M	16M	16M
WORKLOAD	FS7B35R	FS7B35R	FS7B35R	FS7B35R
USERS	240	240	240	240
NORMALIZED BY COMMAND				
Close File Requests	0.3923	0.3907	0.3937	0.3894
Commit Requests	0.0000	0.0000	0.0000	0.0339
Connect Requests	0.0032	0.0043	0.0041	0.0071
Delete File Requests	0.0976	0.0986	0.0979	0.0971
Lock Requests	0.0245	0.0246	0.0246	0.0246
Open File New Requests	0.0016	0.0016	0.0016	0.0016
Open File Read Requests	0.2227	0.2190	0.2245	0.2213
Open File Replace Requests	0.1487	0.1504	0.1473	0.1478
Open File Write Requests	0.0199	0.0196	0.0197	0.0190
Query File Pool Requests	0.0000	0.0001	0.0001	0.0001
Query User Space Requests	0.0222	0.0227	0.0219	0.0220
Read File Requests	0.2275	0.2271	0.2264	0.2213
Refresh Directory Requests	0.0096	0.0094	0.0092	0.0205
Rename Requests	0.0047	0.0049	0.0046	0.0048
Unlock Requests	0.0246	0.0251	0.0246	0.0251
Write File Requests	0.1302	0.1423	0.1390	0.1397
Total File Pool Requests	1.3293	1.3404	1.3393	1.3752
File Pool Request Service Time	113.5989	51.2940	50.3480	48.7258
Local File Pool Requests	1.3293	1.3404	1.3393	1.3752
Begin LUWs	0.4900	0.4870	0.4968	0.5065
Agent Holding Time (msec)	272.8043	276.2660	273.6477	268.4933
SAC Calls	6.0834	6.1047	6.2483	6.2407
Catalog Lock Conflicts	0.0018	0.0007	0.0008	0.0009
Total Lock Conflicts	0.0018	0.0007	0.0008	0.0009
Lock Wait Time (msec)	0.0390	0.0105	0.0151	0.0350
File Blocks Read	0.9361	0.9451	0.9461	0.9301
File Blocks Written	0.5994	0.6270	0.6156	0.6150
Catalog Blocks Read	0.5842	0.5483	0.5642	0.5458
Catalog Blocks Written	0.2284	0.2165	0.2302	0.2143
Control Minidisk Blocks Read	0.0000	0.0000	0.0001	0.0000
Control Minidisk Blocks Written	0.0357	0.0189	0.0216	0.0200
Log Blocks Written	0.5363	0.5584	0.5688	0.5584
Total DASD Block Transfers	2.9201	2.9141	2.9466	2.8837
BIO Requests to Read File Block	0.4879	0.4965	0.4981	0.4894
BIO Requests to Write File Blocks	0.2848	0.3010	0.2955	0.2954
BIO Requests to Read Catalog Blocks	0.5842	0.5483	0.5642	0.5458
BIO Requests to Write Catalog Blocks	0.2284	0.2165	0.2302	0.2143
BIO Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0001	0.0000
BIO Requests to Write Cntrl Mdsk Blks	0.0038	0.0021	0.0024	0.0022
BIO Requests to Write Log Blocks	0.5363	0.5584	0.5688	0.5584
Total BIO Requests	2.1254	2.1228	2.1593	2.1054
Total BIO Request Time (msec)	76.0748	36.3502	36.6819	34.9549

PAGE REFERENCES	91,108	91,108,126	126,174	174
RELEASE	**ESA 1.0**	**ESA 1.1**	**ESA 1.1**	**ESA 1.1**
RUN ID	**H17F0241**	**H14F0241**	**H14F0242**	**H14M0241**
PROCESSOR	**9221-170**	**9221-170**	**9221-170**	**9221-170**
REAL STORAGE	**64M**	**48M**	**48M**	**48M**
EXP. STORAGE	**0M**	**16M**	**16M**	**16M**
WORKLOAD	**FS7B35R**	**FS7B35R**	**FS7B35R**	**FS7B35R**
USERS	**240**	**240**	**240**	**240**
I/O Requests to Read File Blocks	0.4978	0.4612	0.4649	0.4533
I/O Requests to Write File Blocks	0.2948	0.3170	0.3119	0.3123
I/O Requests to Read Catalog Blocks	0.5842	0.5483	0.5642	0.5458
I/O Requests to Write Catalog Blocks	0.2284	0.2165	0.2302	0.2143
I/O Requests to Read Cntrl Mdsk Blks	0.0000	0.0000	0.0001	0.0000
I/O Requests to Write Cntrl Mdsk Blks	0.0053	0.0032	0.0035	0.0033
I/O Requests to Write Log Blocks	0.5363	0.5584	0.5688	0.5584
Total I/O Requests	2.1469	2.1046	2.1436	2.0874
Get Logname Requests	0.0000	0.0031	0.0031	0.0029
Get LUWID Requests	0.0000	0.0031	0.0031	0.0029
Total CRR Requests	0.0000	0.0063	0.0061	0.0058
CRR Request Service Time (msec)	0.0000	0.2098	0.1945	0.2006
Log I/O Requests	0.0000	0.0063	0.0061	0.0058
DERIVED RESULTS				
Agents Held	2.3	2.3	2.3	2.3
Agents In-call	0.9	0.4	0.4	0.4
Avg LUW Time (msec)	556.7	567.3	550.8	530.1
Avg File Pool Request Time (msec)	85.5	38.3	37.6	35.4
Avg Lock Wait Time (msec)	21.7	15.0	18.9	38.9
SAC Calls / FP Request	4.58	4.55	4.67	4.54
Deadlocks (delta)	0	0	0	0
Rollbacks Due to Deadlock (delta)	0	0	0	0
Rollback Requests (delta)	0	0	0	0
LUW Rollbacks (delta)	0	0	0	0
Checkpoints Taken (delta)	19	11	12	11
Checkpoint Duration (sec)	1.6	3.1	3.3	3.1
Seconds Between Checkpoints	192.3	333.3	303.0	333.3
Checkpoint Utilization	0.8	0.9	1.1	0.9
BIO Request Time (msec)	35.79	17.12	16.99	16.60
Blocking Factor (Blocks/BIO)	1.37	1.37	1.36	1.37
Chaining Factor (Blocks/IO)	1.36	1.38	1.37	1.38

Appendix C. Workloads

CMS Intensive (FS7B)

Workload Description

Since the CMS interactive workload needs to reflect the CMS environment, which might include the Shared File System (SFS), the FS7B workload is designed so that it can be used for many types of runs just by changing the search order of the disks/directories accessed as filemodes A through G. It can be an all minidisk workload by accessing all minidisks. It can be a maximum SFS workload by accessing all directories. It can also be a combination of both. This is done with a file called DSKORDER EXEC which resides on the user′s Y-disk and is executed by the PROFILE EXEC on the user′s A-disk. In these three environments, the READ/WRITE data represents end user data and the READ/ONLY data represents system data shared among many users. The following table, which gives the exact virtual machine disk environment, indicates the difference between these cases (FS7B0, FS7B35, FS7BMAX) by listing ′MINIDISK′ for minidisk and ′SFS′ for SFS directory.

FILEMODE	ACCESS	NUMBER OF FILES	FS7B0	FS7B35	FS7BMAX
A	R/W	100	MINIDISK	SFS	SFS
B	R/W	0	MINIDISK	SFS	SFS
C	R/O	500	MINIDISK	MINIDISK	SFS
D	R/W	500	MINIDISK	SFS	SFS
E	R/O	500	MINIDISK	MINIDISK	SFS
F	R/O	500	MINIDISK	MINIDISK	SFS
G	R/O	500	MINIDISK	MINIDISK	SFS
S	R/O	260	MINIDISK	MINIDISK	MINIDISK
Y	R/O	780	MINIDISK	MINIDISK	MINIDISK

NOTES:

- When the C-disk is a minidisk, all of the files on the C-disk have their FSTs saved in a shared segment.
- The HELP disk has the FSTs saved in a shared segment.
- The CMSINST and CMSVMLIB shared segments are used.
- The CMSFILES shared segment is used when SFS is used.
- ALL read-only SFS directories are defined with PUBLIC READ Authority.
- The read/write SFS directory accessed as filemode D is defined with PUBLIC READ and PUBLIC WRITE Authority.
- The read/write SFS directories accessed as filemode A and B are private directories.

FS7B Variations

In addition to testing the three different environments as listed above, there are two possible drivers to run the workload. These drivers are the Full Screen Internal Driver (FSID) and TeleProcessing Network Simulator (TPNS). Each of these six variations has been given a unique identifier as defined below:

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.

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 4.0 (Shared Segments)
HASM	-	Version 2 Release 1.0
PL/I	-	Version 2 Release 1.0

Measurement Methodology

The general methodology is to determine how many users will drive the CPU utilization for the base measurement to 90% and then use this number for all of the measurements made in that environment. For this document, VM/ESA 1.0 GA code was used for the base measurement and it was determined that minidisk-only runs should logon 5860 users and Shared File System (SFS) runs should logon 4800 users.

The FS7B workload uses the Bactrian think time distribution and aims at getting 30 second average think time. The workload also strives to get as much work done as possible during the measurement period. Thus, if the scripts take less time to complete, more scripts will be executed during the measurement period.

Getting a valid measurement takes several steps. First, all of the users are logged on via TPNS (or FSID). The users are then started. Over the course of five minutes, each user selects a script and starts working. This staggering is done so that all of the users do not start scripts at the same time. A stabilization period (typically 45 minutes) is allowed to elapse so that startup anomalies and user synchronization are eliminated. At the conclusion of this period, measurement tools are started simultaneously to acquire measurement data for a 30 minute measurement interval.

FS7B Script Description

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
PRD517	5%	Productivity Aids Session
DCF517	5%	Edit and Script a File
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

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

```
IF A-DISK is a minidisk
THEN
  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 FIXIT 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.
ELSE
  Execute ZELEX 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.
  Set remote on.
END
```

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.
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.
Load the text file (NOMAP option).
Execute the module.
Execute DELEX Exec to clean A-disk.
Query disk, users, and time.

XED117: Edit of a VSBASIC Program

General Description

The script uses XEDIT 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 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.

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.

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 dont issue spool printer.
List and erase output files.
Reset GLOBAL and clear FILEDEFs.

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).
Issue the SENDFILE command.
Send a file to ME (yourself).
Issue the SENDFILE command.
Send a file to ME (yourself).
Issue RDRLIST command, PEEK and DISCARD a file.
Refresh RDRLIST screen, Receive MYID6 EXEC 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.

DCF517: Edit and Script a File

General Description

This script uses XEDIT 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.

PLI317: PL/I Optimizer Session

General Description

XEDIT 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 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 Exploitation

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
IF   TPNS
THEN
  Set autoread on
  Set remote on
END
```

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.
IF FSID
THEN
  Execute PROFLOGF Exec to send SMSG to LOGONSRV and LOGOFF
ELSE
  Logoff user and wait 60 seconds
  Logon user back to original GRAF-ID
  Set autoread on
  Set remote on
END
```

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.

IBM Office Benchmark (IOB V2.1)

Workload Description

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

All of the IOB measurements included in this report use the DisplayWrite/370 2.1.0 and the OfficeVision/VM 1.1.0 Service Level 101 Program Products.

Measurement Methodology

The general methodology was to logon as many users as possible until the processor utilization reached the desired level (typically 90%).

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

Getting a valid measurement takes several steps. First, all of the users are logged on via TPNS (or FSID). The users are then started. Over the course of ten minutes (for 9021-720 runs) or fifteen minutes (for 9021-580 runs), each user selects a script and starts working. This staggering is done so that all of the users do not start scripts at the same time. A stabilization period (typically 30 minutes) is allowed to elapse so that startup anomalies and user synchronization are eliminated. RTM and the internal XXTRANS tool are used to ensure that the system has properly stabilized. At the conclusion of this period, measurement tools are started simultaneously to acquire measurement data for a 30 minute measurement interval.

After the run data is analyzed and looks like it would qualify for IOB certification, the run data is sent to Dallas for certification. All of the runs in this report were certified as valid IOB runs.

IOB Script Descriptions

The IOB workload consists of nine scripts (scenarios). These scripts are listed below with their defined weighting factors:

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

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

Send Note and Process Light Mail

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

Send Note and Process Heavy Mail

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

View Individual Calendar

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

Update Individual Calendar

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

View User Directory

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

Create Small Text Document

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

Revise Small Text Document

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

End/Begin Office

- End or exit the office software program or environment.
- Begin or enter the office software program or environment.

Logon/Logoff System

- Take the option to log off completely from the system.
- Log back onto the system and enter the office environment.

MVS Guest (CB84)

Workload Description

CB84 (Commercial Batch 1984) is a jobstream intended to represent an MVS 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 79 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 IEBGENER 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.

Measurement Methodology

Preliminary CB84 runs are required to ′prime′ the Virtual Lookaside Facility (VLF) with the appropriate modules (i.e., get them loaded into VLF). Data and tuning information from these initial runs are not valid as measurement data.

The measurements begin 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 depends on the workload and on the I/O configuration. It is set, by experimentation, to the number that results in the CCVUTILP AVERAGE of at least 100% during the steady-state portion of the workload execution--this information is in the RMF Trace Activity report. The overall processor utilization must be at least 80%. The batch jobstream is released when the RMF ZZ ACTIVE message appears by using a PF key rather than the $VS command. The RMF data is inspected to ensure I/O balancing and to compute ITR and ETR values to determine the maximum throughput. The system is then loaded with jobs while the queues are held to allow all preliminary work to complete before starting the measurement. The next step is to simultaneously release the queues and start RMF and other measurement tools, e.g., MONITOR, RTM, and an IBM-internal counter program used to automatically stop RMF after the required number of jobs have executed.

The measurements are run twice under the same conditions to validate the results and to show that they are repeatable.

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 five temporary I/O errors over a 10-minute time span
- No abends
- No permanent I/O errors
- No missing I/O interrupts
- Clean EREP

MVS/SP 3.1.0e was used for the MVS guest measurements described in this document.

VSE Guest (PACEX8)

PACE is a synthetic batch workload consisting of 7 unique jobs representative of the commercial environment. As processors became more powerful, PACE was expanded by replicating the 7 jobs, first 4 times for PACEX4 and then 8 times for PACEX8 which is the VSE batch workload used currently.

The seven jobs are as follows:

- Y*n*DL/1
- Y*n*SORT
- Y*n*COBOL
- Y*n*BILL
- Y*n*STOCK
- Y*n*PAY
- Y*n*FORT

There are 8 sets of these jobs used in PACEX8; they are differentiated by the *n* digit in the name (*n* having a value from 1 to 8).

The programs, data, and work space for the jobs are all maintained by VSAM on separate volumes.

The VSE system is configured with the full complement of 12 static partitions (BG, and F1 through FB). F4 through FB are the partitions used to run the workload batch jobs.

The partitions are configured identically except for the job classes. The jobs and the partition job classes are configured so that the jobs are equally distributed over the 8 partitions and so that, at any one time, the jobs currently running are a mixed representation of the 7 jobs.

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

- CICS/ICCF is shut down
- VTAM is shut down
- The LST queue is emptied (PDELETE LST,ALL)

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

As the workload nears completion, various partitions will finish the work allotted to them. The finish time for both the first and last partitions is noted. The difference between these two times should not be more than 1 to 1.5 minutes. If it is more, the jobs and partitions have to be adjusted to get a more even work distribution.

At workload completion, the ITR can be calculated by dividing 56 (the number of batch jobs) by CPU busy time. The CPU busy time is calculated as elapsed (wall clock) time multiplied by CPU busy percent divided by 100.

Measurement data gathered for the VSE guest measurements in this document all used VSE/ESA 1.1.0.

CMS Pipelines

This section describes in detail the test cases and run environments used in the CMS Pipelines performance study.

Virtual Machine Configurations

The following was the search order and configuration of the virtual machine used for the PRPQ 1.1.6 vs. VM/ESA 1.1 REXX/EXECIO/XEDIT vs. CMS Pipelines traces. In addition it was the virtual machine setup for each user in the multi-user benchmark.

```
      VM SIZE: 2M
      VM MODE: XA
CMS BLOCKSIZE: 4K

 Search Order:
 -------------
 U1000    191 linked R/W as A  and contained 100 files
 U1000    111 linked R/W as B  and contained 0 files
 OPERATOR 295 linked R/O as C  and contained 500 files (with shared FST)
 OPERATOR 296 linked R/O as D  and contained 500 files
 OPERATOR 296 linked R/O as E  and contained 500 files
 OPERATOR 296 linked R/O as F  and contained 500 files
 OPERATOR 296 linked R/O as G  and contained 500 files
 U3       191 linked R/W as H  and contained 35 files (with shared FST)
 MAINT    190 linked R/O as S  and contained 260 files
 MAINT    19E linked R/O as Y/S and contained 780 files
```

Shared Segments

CMS Pipelines was not installed in a shared segment for the PRPQ vs. VM traces; however, it was installed in a shared segment for the REXX/EXECIO/XEDIT vs. CMS Pipelines tests. The OPERATOR 295 disk had shared FSTs and CMS was saved for both the traces and the multi-user benchmarks. In addition, the U3 191 disk had shared FSTs for the REXX/EXECIO/XEDIT vs. CMS Pipelines comparisons only.

PRPQ 1.1.6 CMS Pipelines vs. VM/ESA 1.1 CMS Pipelines Commands Traced

The following 6 commands were traced for PRPQ 1.1.6 and VM/ESA 1.1 CMS Pipelines.

1. ***PIPE CMS Q DISK | > QUERY DISK A***
2. ***PIPE < NATHAN NAMES A| CONSOLE***
3. ***PIPE CP Q N|SPLIT ,|STRIP |LOCATE /- DSC/|COUNT LINES|SPEC *-* 1 /Users disconnected/ NEXT| CONSOLE***
4. ***PIPE (end \) < NATHAN NAMES A|c:LOCATE /:nic/|SPEC 24-* 1|JOIN 2 / /| LITERAL Ids:|CONSOLE \c:|SPEC 24-* 1|JOIN 2 / /|LITERAL NAMES:| > NA OUT A***
5. ***PIPE < NATHAN NAMES A|SPEC 1 A|CONSOLE***
6. ***PIPE LITERAL A RECORD | DUP 9 | FANIN | COUNT LINES | CONSOLE***

REXX/EXECIO/XEDIT vs. CMS Pipelines Test Descriptions

The twelve “functions” are described below with the corresponding REXX and CMS Pipelines EXECs.

Test Case 1 - Read in a small file.

REXX1 EXEC

```
/* */
EXECIO * DISKR NATHAN NAMES A (STEM IN.
EXIT
```

PIPE1 EXEC

```
/* */
PIPE < NATHAN NAMES | STEM IN.
EXIT
```

Test Case 2 - Issue CP command storing results.

REXX2 EXEC

```
/* */
EXECIO * CP (STEM IN. STRING Q N
EXIT
```

PIPE2 EXEC

```
/* */
PIPE CP Q N | STEM IN.
EXIT
```

Test Case 3 - Issue CMS command storing results.

REXX3 EXEC

```
/* */
QUERY DISK (FIFO
DO I=1 TO QUEUED()
  PULL IN.I
END I
EXIT
```

PIPE3 EXEC

```
/* */
PIPE CMS QUERY DISK | STEM IN.
EXIT
```

Test Case 4 - Write line of output to a file.

REXX4

```
/* */
EXECIO 1 DISKW OUT OUT A (STRING Hi There
EXIT
```

PIPE4

```
/* */
PIPE LITERAL Hi There | > OUT OUT A
EXIT
```

Test Case 5 - Determine mode of a particular disk in the search order.

REXX5 EXEC

```
/* */
Q DISK (FIFO
mode=
DO QUEUED() WHILE mode=
  PULL label . m1 .
  If label=CMS11 THEN mode=left(m1,1)
END
EXIT
```

PIPE5 EXEC

```
/* */
PIPE CMS Q DISK | LOCATE /CMS11/ | SPEC 13.1 1.1 | VAR MODE
EXIT
```

Test Case 6 - Read in a small file and output to the screen.

REXX6 EXEC

```
/* */
EXECIO * DISKR NATHAN NAMES A (STEM OUT.
DO I=1 TO OUT.0
  SAY OUT.I
END I
EXIT
```

PIPE6 EXEC

```
/* */
PIPE < NATHAN NAMES A | CONSOLE
EXIT
```

Test Case 7 - Issue the CP QUERY NAMES command and determine the number of disconnected users

REXX7 EXEC

```
/* */
EXECIO * CP (STEM OUT. STRING QUERY NAMES
NUMUSERS=0
K=0
DO I=1 TO OUT.0
  DO UNTIL K=0
    K=POS(-DSC,OUT.I,K+1)
    IF K/=0 THEN
       NUMUSERS=NUMUSERS+1
  END
END I
SAY NUMUSERS users disconnected
EXIT
```

PIPE7 EXEC

```
/* */
PIPE CP Q N|SPLIT ,|LOCATE /-DSC/|COUNT LINES|,
   SPEC *-* 1 /users disconnected/ NEXT | CONSOLE
EXIT
```

Test Case 8 - Read in a file and output certain lines to the screen (matching a search criteria) while writing the rest to disk.

REXX8 EXEC

```
/* */
I=0
J=0
SAY Ids:
EXECIO 1 DISKW NA OUT A   (STRING NAMES:
OUT=
OUT1=
EXECIO 1 DISKR NATHAN NAMES A
DO WHILE RC=0
  PARSE PULL LINE 1 FLAG . 24 REST
  IF LEFT(FLAG,4)=:nic THEN DO
    I=I+1
    OUT=OUT||REST||
```

```
    IF I=3 THEN DO
      SAY OUT
      OUT=
      I=0
    END
  END
  ELSE DO
    J=J+1
    OUT1=OUT1||REST||
    IF J=3 THEN DO
      EXECIO 1 DISKW NA OUT A (VAR OUT1
      OUT1=
      J=0
    END
  END
  EXECIO 1 DISKR NATHAN NAMES A
END
IF I>0 THEN SAY OUT
IF J>0 THEN EXECIO 1 DISKW NA OUT A (VAR OUT1
FINIS NA OUT A
FINIS NATHAN NAMES A
EXIT
```

REXX8B EXEC

```
/* */
I=2; I1=0
J=2; J1=0
OUT.1 =Ids:
OUT1.1=NAMES:
OUT.2=
OUT1.2=
EXECIO * DISKR NATHAN NAMES A (STEM IN.
DO K=1 TO IN.0
  PARSE VAR  IN.K 1 FLAG . 24 REST
  IF LEFT(FLAG,4)=:nic THEN DO
    I1=I1+1
    OUT.I=OUT.I||REST||
    IF I1=3 THEN DO
      I=I+1
      OUT.I=
      I1=0
    END
  END
  ELSE DO
    J1=J1+1
    OUT1.J=OUT1.J||REST||
    IF J1=3 THEN DO
      J=J+1
      OUT1.J=
      J1=0
    END
  END
END
IF I1=0 THEN I=I-1
IF J1=0 THEN J=J-1
DO K= 1 TO I
  SAY OUT.K
END K
EXECIO J DISKW NA OUT A (STEM OUT1.
FINIS NA OUT A
FINIS NATHAN NAMES A
EXIT
```

PIPE8 EXEC

```
/* */
PIPE (END \) < NATHAN NAMES A|C:LOCATE /:nic/|SPEC 24-* 1|JOIN 2 //|,
 LITERAL Ids:|CONS \C:|SPEC 24-* 1|JOIN 2 //|LITERAL NAMES:| > NA OUT A
EXIT
```

Test Case 9 - Read in a selected portion of a file, storing results.

REXX9 EXEC
```
/* */
QUEUE FINDIT
QUEUE QQ
XEDIT RUN MASTER H (NOPROF
EXIT
```

FINDIT XEDIT
```
/* */
FIND L23F1481
IF RC=0 THEN DO
 EXTRACT /CURLINE
  I=0
  DO WHILE LEFT(CURLINE.3,8) /=--------
    I=I+1
    ENTRY.I=CURLINE.3
   NEXT
   EXTRACT /CURLINE
  END
  I=I+1
  ENTRY.0=I
  ENTRY.I=CURLINE.3
END
EXIT
```

REXX9B EXEC
```
/* */
EXECIO * DISKR RUN MASTER H (FI /L23F1481/ STEM ENTRY.
I=1
DO WHILE LEFT(ENTRY,8)/=--------
  I=I+1
 EXECIO 1 DISKR RUN MASTER K (VAR ENTRY
  ENTRY.I=ENTRY
END
EXIT
```

PIPE9 EXEC
```
/* */
PIPE < RUN MASTER H | BETWEEN /L23F1481/ /--------/ | STEM ENTRY.
EXIT
```

Test Case 10 - Read in and output to the screen lines of a file matching a certain search criteria.

REXX10 EXEC
```
/* */
XEDIT SFSDUMP CONTROL H (PROF BETWIXT
EXIT
```

BETWIXT XEDIT
```
/* */
TOP
ALL /AREA/
EXTRACT /CURLINE
DO while curline.3/=
  PARSE VALUE CURLINE.3 with . = out , .
  SAY out
  NEXT
  EXTRACT /CURLINE
END
QQ
EXIT
```

PIPE10 EXEC
```
/* */
PIPE < SFSDUMP CONTROL H|FIND AREA|BETWIXT|CONS
EXIT
```

BETWIXT REXX
```
/* */
PARSE Arg n1 n2
READTO In
 DO while rc=0
```

```
   PARSE VALUE In with .= out , .
   Output OUT
   READTO In
 END
EXIT
```

PIPE10B EXEC

```
/**/
PIPE < SFSDUMP CONTROL H|BETWIXT2|FIND  |CONS
EXIT
```

BETWIXT2 REXX

```
/**/
READTO In
 DO while rc=0
   PARSE VALUE In with . AREA= out , .
   Output OUT
   READTO In
 END
EXIT
```

Test Case 11 - Read a reader spool file into storage.

REXX11 EXEC

```
/* */
EXECIO * CARD (STEM IN.
EXIT
```

PIPE11 EXEC

```
/**/
PIPE READER | DROP 1 | SPEC 2-* 1 | STEM IN.
EXIT
```

Test Case 12 - Read in a reader spool file, reorganize columns of data and output to disk.

REXX12 EXEC

```
/* */
EXECIO * CARD (STEM IN.
DO I=1 TO IN.0
  OUT.I=substr(IN.I,19,8)||          ||substr(IN.I,1,8)
END I
EXECIO IN.0 DISKW OUT MAP A 1 F 40 (STEM OUT. FINIS
EXIT
```

PIPE12 EXEC

```
/**/
PIPE READER | DROP 1 | SPEC 20-27 1 2-9 19-40 |  > OUT MAP A FIXED 40
EXIT
```

INSTVER Communications Benchmark

The INSTVER benchmark is the VM PWSCS installation verification program. INSTVER is coded in the C programming language and conforms to the Common Programming Interface for Communications (CPI-C) architecture. The INSTVER benchmark consists of two logically distinct programs-- a client program and a server program. The criterion that is measured by these two programs is data throughput. To ensure that the measurements attained while running INSTVER were as accurate as possible, system activity was quiesced while the two benchmark programs were running.

The INSTVER benchmark is started by invoking the client program. The client program processes the command line arguments which include the message size, iteration count and symbolic destination name. The message size determines the amount of data the benchmark will send on each of the sequential sends and receives. The iteration count is used to smooth out any irregularities that may be introduced due to the resolution of the system clock. The symbolic destination name is used by CPI-C to determine where the server resource is located, what transaction program on the remote server needs to be invoked and, if conversation security is being used, the userid and password of the server machine.

The client initializes and allocates a conversation to the resource specified by the symbolic destination name. As a result of the allocation, the server process begins execution on the remote system. The INSTVER benchmark uses sync-level processing to control the data flow, as well as the timings in the benchmark. When the server confirms the conversation allocation, the client sends a message to the server which contains the size and number of messages that the server is to receive and send. When the server receives this message it issues a confirm, which lets the client know that the benchmark portion of code is ready to commence.

The client takes an initial reading of the system clock and then streams messages of the appropriate size and number. The server sits in a loop receiving the messages. On the last send, the client requests confirmation that all of the data was received. When the confirmation is received by the client, another reading of the clock is taken and the data throughput is calculated. The direction of the conversation is then changed, with the server sending data to the client. Calculations are then made for the throughput of the returning data stream. The overall throughput is the aggregate throughput for both halves of the benchmark.

When considering communications benchmarks, it is important to consider the data flow that is employed. The two most prevalent flows that are used are wrapping and streaming. In a wrap data flow, each iteration consists of a send to the server followed by a receive by the client. In a streaming data flow, the client issues several sends, one after the other. The INSTVER benchmark uses a streaming data flow.

The iteration count is employed to ensure that system timings are as accurate as possible. The resolution of the system clock can be a large factor in the benchmark if few iterations are used. For example, if one iteration takes 5 milliseconds, then a 1 millisecond error in measuring the system clock will result in a 20 percent error in measured throughput. If several iterations are run between the clock readings, then the overall clock error will be much less than the overall time of the benchmark. Thus, the timer errors tend to get cancelled out the longer the benchmark runs. If the benchmark using the iteration count takes 20 seconds, the percentage error due to the system clock is extremely small.

The reason wrap data flows are usually slow in comparison to streaming flows is that in the CPI-C architecture the conversation state must change each time there is a change of direction in data flow. CPI-C is a half-duplex protocol, meaning that data flows over the wire in one direction only at any given instant in time. ISFC implements a pacing algorithm in order to balance out the workload on the server and client. Pacing is the number of messages that a client can send to a server before a pacing response is required. This mechanism prevents a client from swamping a server with more data than can reasonably be handled.

The streaming data flow was chosen for its similarity to many common communications transactions (such as file transfer) as well as its ease of implementation. Ease of implementation was also enhanced by writing the host server in CPI-C as opposed to using the assembler language APPC interface. CPI-C is the IBM Systems Application Architecture (SAA) standard protocol for communications programs. It is highly portable, supporting a wide variety of systems. In addition, it provides many capabilities such as sync-level processing, security verification, and directory management. The API provided by the VM PWSCS software used on the LAN conforms to the CPI-C architecture. Writing the host server application in CPI-C preserves symmetry with the code running on the client workstation.

On System/390 systems, CPI-C programs are converted to the underlying APPC protocols seamlessly. Choosing to write a server that is portable and symmetric to the client means that host performance is sacrificed somewhat. This sacrifice is justified by the desire to measure the new function as it would most commonly be used in a production environment. It is thought to be a reasonable assumption that programmers using the CPI-C interface in VM PWSCS on the workstation will write their server application to the same interface on the host. The abundance of processing power on the System/390 helps to lessen the need to write communications programs in assembler using APPC. The INSTVER benchmark is an attempt to get throughput figures for a Communication Services collection in a typical configuration using the standard SAA-approved CPI-C protocol.

Appendix D. Configuration Details

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*
9021-720	227576
9021-580	227576
9021-580	229910
9121-480	C23074
9121-320	C23074
9121-320	C23070
9221-170	95D + MC85871.B22 patch

Named Saved Segments / Systems

CMS allows the use of saved segments for shared code. Using saved segments can greatly improve performance by reducing end users′ working set sizes and thereby decreasing paging. The environments in this report use the following saved segments:

- CMS: Contains the CMS nucleus and File Status Tables for the S and Y disks.
- CMSFILES: Contains the SFS server code modules DMSDAC and DMSSAC.
- CMSINST: Contains the EXECs-in-storage segment.
- CMSVMLIB: Contains the CSL code.
- HELP: Contains FSTs for the HELP disk.
- GOODSEG: Contains FSTs for the C disk.
- FORTRAN: This segment space has 2 members: DSSVFORT for the FORTRAN compiler and FTNLIB10 for the Library composite modules.
- DSMSEG3: Contains Document Composition Facility (DCF).
- OFSSEG: Contains OV/VM user functions.
- DW370: Contains the DW370 module.
- DDDCL210: Contains the DW370 compiled CLISTS.
- DW362: Contains FSTs for the DW/370 362 disk.
- ADM399: Contains FSTs for the OV/VM 399 disk.
- GCSXA: Contains the GCS nucleus.
- VTAMXAA: Contains the VTAM code.

Server Options

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 OPERATOR MBUILD BUILD U1 U2 U3 U4 U5
FILEPOOLID FPn
NOBACKUP
FORMAT
USERS nnn
FULLDUMP
SAVESEGID
MSGS
ACCOUNT
```

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. The table that follows contains the USERS settings for the set of processors reported on in this document.

PROCESSOR	*USERS*
9021-720	1200
9021-580	1500
9121-480	1210
9121-320	1210
9221-170	240

For more information on SFS and SFS tuning parameters, check the *VM/ESA 1.1 CMS Planning and Administration* manual.

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 OPERATOR MBUILD BUILD U1 U2 U3 U4 U5
NOBACKUP
FULLDUMP
SAVESEGID
FILEPOOLID CRRFPA
ACCOUNT
NOFORMAT
CRR
LUNAME nodeid.userid
USERS 500
MSGS
```

Appendix E. VTAM Definition Statements

This section describes the VTAM definition statements used for the measurement runs in this report. By referring to the detailed configuration sections in Part 3, "Specific Measurements" on page 47, the VTAM definition statements used for each particular measurement can be determined as follows:

1. For runs made on the 9221 processor, refer to "VTAM V3R3 / 9221-170 / TPNS Driver via CTCA" on page 306.
2. For runs which used VM/XA 2.1 and VTAM V3R2, refer to "VTAM V3R2 / 9021-720 / TPNS Driver via 3745s." There were only two runs, listed in "OfficeVision Migration from VM/XA 2.1" on page 129, which used this configuration.
3. For runs which list a 3745 as the communications controller, and/or list VTAMXAA, VSCSXA2, or VSCSXA3 as server machines, refer to "VTAM V3R3 / 9021-720 / 9121-480 / TPNS Driver via 3745s" on page 304.
4. For runs which list a 3088 as the communications controller, and/or list VTAM as a server machine, refer to "VTAM V3R3 / 9021-580 / TPNS Driver via CTCA" on page 305.

VTAM V3R2 / 9021-720 / TPNS Driver via 3745s

This section describes the VTAM V3R2 definition statements used for the VM/XA 2.1 measurement runs made on a 9021-720, and driven by a TPNS system via channel-attached 3745s. There were three VTAM machines (VTAMA, VTAMB, and VTAMC), each with an internal VSCS, and three external VSCS machines (VSCS2, VSCS4, and VSCS6).

VTAM Machine - VTAMA

This section describes the VTAM definition statements which are specific to this VTAM machine (and internal VSCS).

VTAM Start Options - ATCSTR45

```
HOSTSA=06,
PROMPT,
NOTRACE,
TYPE=VTAM,
MAXSUBA=31,
NETID=NET1,
CONFIG=45,
SSCPID=06,
SSCPNAME=TEST,
IOBUF=(1000,256,19,,50,50),
CRPLBUF=(500,,15,,80,80),
LFBUF=(25,,0,,10,1),
LPBUF=(100,,15,,50,50),
SFBUF=(80,,0,,50,1),
WPBUF=(3070,,0,,10,1)
```

Configuration List - ATCCON45

```
VTAMAPPL
```

Application Major Node - VTAMAPPL

```
VTAMAPPL VBUILD TYPE=APPL
VMBUB    APPL   AUTH=(BLOCK,PASS,ACQ),PARSESS=YES,ACBNAME=VMBUB,
                AUTHEXIT=YES
VSCS2    APPL   AUTH=(BLOCK,PASS,ACQ),PARSESS=YES,ACBNAME=VSCS2,
                AUTHEXIT=YES
```

VSCS Start Options - DTIUSER1

```
VMBUB   DTIGEN DTIUSER=1,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
               APPLID=VMBUB,VSAMLM=8,RPLNUM=16
```

VTAM Machine - VTAMB

This section describes the VTAM definition statements which are specific to this VTAM machine (and internal VSCS).

VTAM Start Options - ATCSTR45

```
HOSTSA=06,
PROMPT,
NOTRACE,
TYPE=VTAM,
MAXSUBA=31,
NETID=NET1,
CONFIG=45,
SSCPID=06,
SSCPNAME=TEST,
IOBUF=(1000,256,19,,50,50),
CRPLBUF=(500,,15,,80,80),
LFBUF=(25,,0,,10,1),
LPBUF=(100,,15,,50,50),
SFBUF=(80,,0,,50,1),
WPBUF=(3070,,0,,10,1)
```

Configuration List - ATCCON45

```
VTAMAPPL
```

Application Major Node - VTAMAPPL

```
VTAMAPPL VBUILD TYPE=APPL
VSCS3    APPL   AUTH=(BLOCK,PASS,ACQ),PARSESS=YES,ACBNAME=VSCS3,
                AUTHEXIT=YES
VSCS4    APPL   AUTH=(BLOCK,PASS,ACQ),PARSESS=YES,ACBNAME=VSCS4,
                AUTHEXIT=YES
```

VSCS Start Options - DTIUSER3

```
VSCS3  DTIGEN DTIUSER=3,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
              APPLID=VSCS3,VSAMLM=8,RPLNUM=16
```

VTAM Machine - VTAMC

This section describes the VTAM definition statements which are specific to this VTAM machine (and internal VSCS).

VTAM Start Options - ATCSTR45

```
HOSTSA=06,
PROMPT,
NOTRACE,
TYPE=VTAM,
MAXSUBA=31,
NETID=NET1,
CONFIG=45,
SSCPID=06,
SSCPNAME=TEST,
IOBUF=(1000,256,19,,50,50),
CRPLBUF=(500,,15,,80,80),
LFBUF=(25,,0,,10,1),
LPBUF=(100,,15,,50,50),
SFBUF=(80,,0,,50,1),
WPBUF=(3070,,0,,10,1)
```

Configuration List - ATCCON45

```
VTAMAPPL
```

Application Major Node - VTAMAPPL

```
VTAMAPPL VBUILD TYPE=APPL
VSCS5    APPL   AUTH=(BLOCK,PASS,ACQ),PARSESS=YES,ACBNAME=VSCS5,
                AUTHEXIT=YES
VSCS6    APPL   AUTH=(BLOCK,PASS,ACQ),PARSESS=YES,ACBNAME=VSCS6,
                AUTHEXIT=YES
```

VSCS Start Options - DTIUSER5

```
VSCS5  DTIGEN DTIUSER=5,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
              APPLID=VSCS5,VSAMLM=8,RPLNUM=16
```

External VSCS Machine - VSCS2

This section describes the VTAM definition statements which are specific to this external VSCS machine.

VSCS Start Options - DTIUSER2

```
VSCS2  DTIGEN DTIUSER=2,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
             APPLID=VSCS2,VSAMLM=8,RPLNUM=16
```

External VSCS Machine - VSCS4

This section describes the VTAM definition statements which are specific to this external VSCS machine.

VSCS Start Options - DTIUSER4

```
VSCS4  DTIGEN DTIUSER=4,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
             APPLID=VSCS4,VSAMLM=8,RPLNUM=16
```

External VSCS Machine - VSCS6

This section describes the VTAM definition statements which are specific to this external VSCS machine.

VSCS Start Options - DTIUSER6

```
VSCS6   DTIGEN DTIUSER=6,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
               APPLID=VSCS6,VSAMLM=8,RPLNUM=16
```

VTAM V3R3 / 9021-720 / 9121-480 / TPNS Driver via 3745s

This section describes the VTAM V3R3 definition statements used for the VM/ESA 1.0 and VM/ESA 1.1 measurement runs made on a 9021-720 and 9121-480, and driven by a TPNS system via channel-attached 3745s. There was one VTAM machine (VTAMXAA), with an internal VSCS, on both systems. The 9021-720 had two external VSCS machines (VSCSXA2 and VSCSXA3), and the 9121-480 had one external VSCS machine (VSCSXA2).

VTAM Machine - VTAMXAA

This section describes the VTAM definition statements which are specific to this VTAM machine (and internal VSCS).

VTAM Start Options - 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)
```

Configuration List - ATCCON45

```
VTAMAPPL
```

Application Major Node - VTAMAPPL

```
VTAMAPPL VBUILD TYPE=APPL
VMBUB    APPL  AUTH=(PASS,ACQ),PARSESS=YES,ACBNAME=VMBUB,
               AUTHEXIT=YES
VSCSXA2  APPL  AUTH=(PASS,ACQ),PARSESS=YES,ACBNAME=VSCSXA2,
               AUTHEXIT=YES
VSCSXA3  APPL  AUTH=(PASS,ACQ),PARSESS=YES,ACBNAME=VSCSXA3,
               AUTHEXIT=YES
```

VSCS Start Options - DTIUSER1

```
VMBUB    DTIGEN DTIUSER=1,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
               APPLID=VMBUB,VSAMLM=8,RPLNUM=16
```

External VSCS Machine - VSCSXA2

This section describes the VTAM definition statements which are specific to this external VSCS machine.

VSCS Start Options - DTIUSER2

```
VSCSXA2  DTIGEN DTIUSER=2,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
               APPLID=VSCSXA2,VSAMLM=8,RPLNUM=16
```

External VSCS Machine - VSCSXA3

This section describes the VTAM definition statements which are specific to this external VSCS machine.

VSCS Start Options - DTIUSER3

```
VSCSXA3  DTIGEN DTIUSER=3,SPEC=N,SCHED=Y,SCIPCNT=1,FSREAD=N,
               APPLID=VSCSXA3,VSAMLM=8,RPLNUM=16
```

VTAM V3R3 / 9021-580 / TPNS Driver via CTCA

This section describes the VTAM V3R3 definition statements used for the VM/ESA 1.0 and VM/ESA 1.1 measurement runs made on a 9021-580, and driven by a TPNS system via a channel-to-channel adapter. There was one VTAM machine (VTAM), with an internal VSCS.

VTAM Machine - VTAM

This section describes the VTAM definition statements which are specific to this VTAM machine (and internal VSCS).

VTAM Start Options - ATCSTRA1

```
SSCPID=5392,
MAXSUBA=40,
CONFIG=A1,
HOSTSA=01,
PROMPT,
DLRTCB=4,
SUPP=NOSUP,
NETID=NETSNA,
SSCPNAME=BIG,
NOTRACE,
TYPE=VTAM,
LPBUF=(150,,15,,50,50),
LFBUF=(25,,0,,10,1),
WPBUF=(7710,,00,,10,1),
SFBUF=(80,,0,,50,1),
CRPLBUF=(3550,,15,,80,80),
IOBUF=(1500,256,19,,50,50)
```

Configuration List - ATCCONA1

```
ISTAPPLS,PATH0102,CTCBIG,CDRMSBIG
```

Application Major Node - ISTAPPLS

```
VM       APPL  AUTH=(PASS,ACQ),ACBNAME=VM,PRTCT=VM,AUTHEXIT=YES,
               SONSCIP=YES
```

Path Table - PATH0102

```
PATH12   PATH  DESTSA=02,
               ER0=(02,1),ER1=(02,1),
               VR0=0,VR1=1
```

Channel-Attachment Major Node - CTCBIG

```
CTCABIG  VBUILD TYPE=CA
CTCBIGG  GROUP  LNCTL=CTCA,ISTATUS=ACTIVE,DELAY=0.1,REPLYTO=25
CTCBIGL  LINE   ADDRESS=5E3,ISTATUS=ACTIVE,MAXBFRU=(10,32)
CTCBIGP  PU     ISTATUS=ACTIVE
```

CDRM Major Node - CDRMSBIG

```
CDRMS    VBUILD TYPE=CDRM
CDRM12   CDRM  SUBAREA=01,ISTATUS=ACTIVE,CDRSC=OPT,CDRDYN=YES
CDRM21   CDRM  SUBAREA=02,ISTATUS=ACTIVE,CDRSC=OPT,CDRDYN=YES
```

VTAM V3R3 / 9221-170 / TPNS Driver via CTCA

This section describes the VTAM V3R3 definition statements used for the VM/ESA 1.0 and VM/ESA 1.1 measurement runs made on a 9221-170, and driven by a TPNS system via a channel-to-channel adapter. There was one VTAM machine (VTAM), with an internal VSCS.

VTAM Machine - VTAM

This section describes the VTAM definition statements which are specific to this VTAM machine (and internal VSCS).

VTAM Start Options - ATCSTRH3

```
SSCPID=5392,
MAXSUBA=40,
CONFIG=H3,
HOSTSA=01,
PROMPT,
DLRTCB=4,
SUPP=NOSUP,
NETID=NETSNA,
SSCPNAME=BIG,
NOTRACE,
TYPE=VTAM,
LPBUF=(150,,15,,50,50),
LFBUF=(25,,0,,10,1),
WPBUF=(1550,,00,,10,1),
SFBUF=(80,,0,,50,1),
CRPLBUF=(1550,,15,,80,80),
IOBUF=(1500,256,19,,50,50)
```

Configuration List - ATCCONH3

```
ISTAPPLS,PATH0102,CTC5F3,CDRMSBIG
```

Application Major Node - ISTAPPLS

```
VM       APPL  AUTH=(PASS,ACQ),ACBNAME=VM,PRTCT=VM,AUTHEXIT=YES,
               SONSCIP=YES
```

Path Table - PATH0102

```
PATH12   PATH  DESTSA=02,
               ER0=(02,1),ER1=(02,1),
               VR0=0,VR1=1
```

Channel-Attachment Major Node - CTC5F3

```
CTCA5F3  VBUILD TYPE=CA
CTC5F3G  GROUP  LNCTL=CTCA,ISTATUS=ACTIVE,DELAY=0.2,REPLYTO=25
CTC5F3L  LINE   ADDRESS=5F3,ISTATUS=ACTIVE,MAXBFRU=(10,32)
CTC5F3P  PU     ISTATUS=ACTIVE
```

CDRM Major Node - CDRMSBIG

```
CDRMS    VBUILD TYPE=CDRM
CDRM12   CDRM  SUBAREA=01,ISTATUS=ACTIVE,CDRSC=OPT,CDRDYN=YES
CDRM21   CDRM  SUBAREA=02,ISTATUS=ACTIVE,CDRSC=OPT,CDRDYN=YES
```

Appendix F. Customer Performance Experience

Overview

This section describes how the migration to VM/ESA Release 1.1 affected the performance of one installation that participated in the Early Support Program (ESP) for VM/ESA Release 1.1.[4]

This customer is a large public utility. This datacenter (which runs both a VM system and an MVS system) acts as a service bureau, supplying computing services to over 5,000 users throughout the corporation in the areas of financial, customer, business and engineering applications. Primary applications on the VM system are office automation (OV/VM) and database (SQL/DS). There are 5 major SQL servers utilizing approximately 35 GB of DASD and one-third of the CPU resources at peak periods. This system runs on an ES/3090 processor model 300J, and was running VM/XA SP Release 2.1. prior to the migration to VM/ESA Release 1.1.

Presented first is a summary of the impact on performance that this customer saw as a result of the migration to VM/ESA Release 1.1. The performance characteristics of this system will then be examined in more detail, both before and after the migration. The analysis of the performance data was done using VM Performance Planning Facility (VMPPF). VMPPF is a performance-management and capacity-planning tool which combines a powerful modeling facility with extensive data-reduction and workload-classification capabilities.

Note: The performance data gathered from this customer differs from the data presented elsewhere in this publication in the following ways:

- The performance data presented in the body of this publication was collected in a laboratory environment, with a controlled workload, while the data collected from this customer represents a real production environment, with a variable, uncontrolled workload.
- The response times reported in the laboratory measurements are external response times, while the response times shown for this customer are internal response times reported by VMPPF.
- During the period in which performance data was collected, this customer migrated only CP. The migration from CMS 5.6 to CMS 8 was completed after the last sample of performance data was collected. Therefore, the performance numbers shown here reflect only the changes introduced by CP and not those introduced by CMS 8.

[4] This customer's experience should not be viewed as typical of what the majority of customers will experience, nor should it be viewed as typical of what the ESP customers experienced. Some saw better results, some saw worse. It was selected for inclusion in this document because it was one case where the performance data collected showed a definite change as a result of the migration to VM/ESA Release 1.1.

Summary of Performance Changes

Prior to the migration, this system was supporting 1445 logged-on users on an ES/3090 processor model 300J, with 128 MB of central storage and 128 MB of expanded storage. The processor utilization was 76% (54% Problem, 22% CP) and channel utilization never exceeded 16% on any channel. Storage Utilization (as calculated by VMPPF)[5] was 27%. Trivial response time was less than 0.4 seconds.

When using VMPPF to analyze system performance, a variable that is often used to measure the work being done by the system is the number of non-idle users.[6] Prior to the migration the number of non-idle users was 1020.[7]

As the migration to VM/ESA Release 1.1 progressed, the workload on the system (as measured by non-idle users) increased by 26%, from 1020 to 1288 non-idle users. Because of the migration to VM/ESA Release 1.1, the system was able to absorb this growth with no change required to its existing configuration. There was only about a 12% increase in processor utilization. Overall, trivial response time, which was excellent from the outset, increased slightly (probably imperceptibly). And, in general, non-trivial response time improved slightly.

This was possible in large measure because of increased efficiency in VM/ESA code and the significant increase in dynamic pageable space (24 MB) due to VM/ESA′s use of pageable Page Management Blocks.[8]

Methodology

On three separate occasions within a space of approximately six months, VM monitor data was gathered. Representative peak-hour samples were chosen for in-depth analysis from each set of data. The charts which follow are black-and-white renderings of the multi-color graphics produced by VMPPF.

The first sample was taken prior to the VM/ESA Release 1.1 migration. The second sample, approximately four months later, was taken soon after the migration was complete. The third sample was taken approximately two-months later. Having this third sample point is very handy. It serves as a point of reference for data validation and growth projections.

[5] VMPPF defines Storage Utilization as the percentage of primary pageable storage utilized by the in-queue users.

[6] VMPPF classifies a user as ***non-idle*** if he records almost ***any*** activity during the measured period. VMPRF, on the other hand, counts users during each monitor interval (usually 1 minute). A user must record activity during a monitor interval to be considered as ***active*** during that monitor interval.

[7] By comparison, according to VMPRF, there was an average of 256 active users.

[8] In VM/XA and VM/ESA, there is a Page Management Block (PGMBK) for each 1 MB segment of a user′s virtual storage. The PGMBK contains the page table, the page status table, and the auxiliary storage address table. In VM/XA, these blocks are fixed in real storage. In VM/ESA (beginning with Release 1.0), these control blocks are pageable. In VM/SP HPO, the equivalent control blocks (page and swap tables) are also pageable.

Original Customer Environment

VMPPF provides for the analysis of the data by workload distribution. The I/O rate (virtual I/O, or VIO) was chosen as the basis for categorizing this workload. Major service machine applications are identified separately. Note that over 90% of the users are consuming only about 40% of the CPU resources. This is a typical resource-consumption pattern for CMS-intensive workloads.

CLASS	VIO RATE PER HOUR	NUMERICAL DISTRIBUTION	CPU-USAGE DISTRIBUTION
VIO0	0	6.8%	0.4%
VIO1	< 1,000	75.9%	29.1%
VIO2	< 2,500	10.6%	11.0%
VIO3	< 5,000	2.8%	5.1%
VIO4	< 20,000	1.2%	3.0%
VIO5	< 25,000	0.7%	11.5%
VIO6	< 50,000	0.2%	0.4%
VIO7	< 100,000	0.1%	1.3%
RSCS		0.1%	0.7%
OV-SVM		0.6%	0.5%
SQL-SVM		1.1%	7.6%
VTAM		0.1%	4.8%
		100.2%	75.4%

- VIO0 consists mostly of special-purpose servers and "nearly idle" users;
- VIO1 and VIO2 are the typical CMS users; note that the vast majority of users fall into one of these two classes;
- VIO3 through VIO5 are "power" users and heavy database users;
- VIO6 and VIO7 are special-purpose server machines.
- OV-SVM is an aggregate of the OVVM calendar, database, and mailbox server machines.
- SQL-SVM is an aggregate of the various SQL/DS server machines on the system.

Figure 11 is a visual breakdown of the components of response time for each workload class. The first column of numbers lists response times corresponding to trivial and non-trivial transactions for each user class. The class name is identified to the immediate right. The center of the chart displays the components of each transaction type, broken into service and wait times for CPU, I/O, paging, external time, and time waiting for identified server virtual machines.

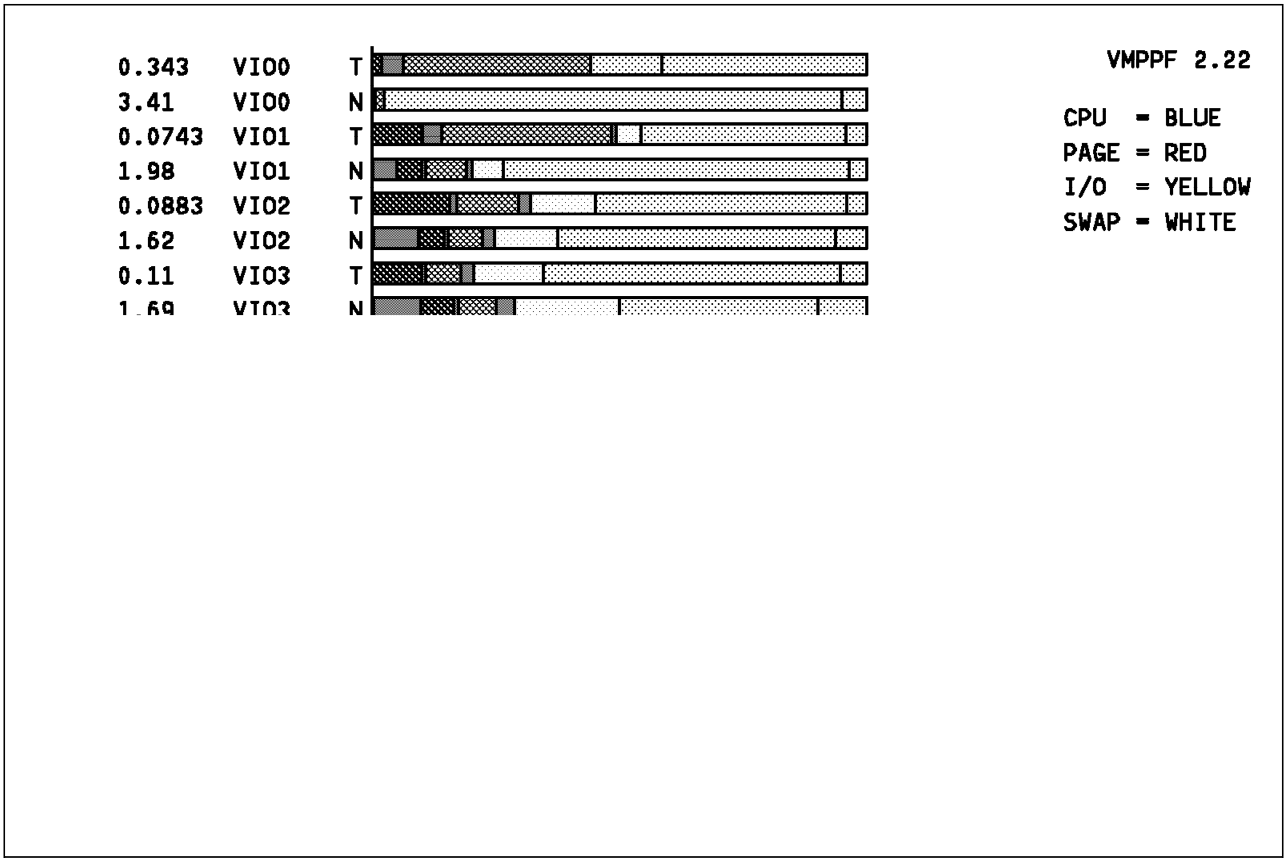

Figure 11. Transaction Profile Summary

Note that there are significant differences in the response-time components between classes, and even within a class between trivial and non-trivial transactions. This illustrates why different types of workloads will see different results when any improvements (including an operating system upgrade such as this) are made to the system.

Trivial response time for all classes averages far less than a half second. It would appear that interactive performance is very satisfactory.

VTAM's transaction (referred to as "***work unit***") rate is over 350,000 per hour. Yet each individual work unit takes practically no time at all (2.8 ***milliseconds***). This indicates that service to remote terminals should also be excellent.