International Technical Support Organization

SG24-4674-00

VSE/ESA Version 2 Release 1
The Turbo Dispatcher

February 1996

International Technical Support Organization
Boeblingen Center

International Technical Support Organization

SG24-4674-00

VSE/ESA Version 2 Release 1 The Turbo Dispatcher

February 1996

Take Note!

Before using this information and the product it supports, be sure to read the general information under "Special Notices" on page xi.

First Edition (February 1996)

This edition applies to Virtual Storage Extended/Enterprise System Architecture V2R.1.1, program number 5690-VSE.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

An ITSO Technical Bulletin Evaluation Form for reader′s feedback appears facing Chapter 1. If the form has been removed, comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. 3222 Building 71032-02
Postfach 1380
71032 Boeblingen, Germany

Abstract

This document highlights selected aspects of the VSE/ESA Turbo Dispatcher introduced in VSE/ESA 2.1. Various VSE/ESA 2.1 facilities were tested and are discussed in detail, to illustrate how the VSE/ESA 2.1 user can effectively exploit the Turbo Dispatcher to improve the performance of a given workload in a VSE/ESA 2.1 multiprocessor environment.

Guidelines for using the VSE/ESA 2.1 Turbo Dispatcher are provided as well as examples of the system setup and the results of the different test runs.

This manual is written mainly for customer and IBM technical personnel responsible for the planning and implementation of VSE/ESA 2.1 multiprocessor support in their environments.

Basic knowledge of VSE/ESA 2.1 and VSE/ESA tuning is assumed.

(58 pages)

Contents

Figures

Tables

Special Notices

This publication is intended to help customer and IBM technical personnel to:

- Understand better the concepts and performance aspects of the VSE/ESA 2.1 Turbo Dispatcher
- Improve the utilization of a multiprocessor system in a VSE/ESA 2.1 environment

The information in this publication is not intended as a specification of any programming interfaces that are provided by VSE/ESA 2.1. See the Publications section of the IBM Programming Announcement for VSE/ESA 2.1, program number 5690-VSE, for more information about what publications are considered to be product documentation.

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 product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. 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.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ACF/VTAM	CICS
CICS/VSE	DB2
ECKD	ES/9000
ESA/370	ESA/390
IBM	IMS
PR/SM	S/390
VM/ESA	VSE/ESA
VTAM	

The following terms are trademarks of other companies:

C-bus	Corollary, Inc.
PC Direct	Ziff Communications Company (used by IBM Corporation under license)
UNIX	X/Open Company Ltd. (registered trademark)
Windows	Microsoft Corporation

Other trademarks are trademarks of their respective companies.

Preface

The purpose of this document is to help the reader better understand the concepts and performance aspects of the VSE/ESA 2.1 Turbo Dispatcher. It contains an explanation of the Turbo Dispatcher concepts, and detailed performance tests of various system facilities in a VSE/ESA 2.1 multiprocessor environment.

This document is intended for customer and IBM technical personnel who have the responsibility to plan and implement the VSE/ESA 2.1 multiprocessor support in their environments.

How This Document is Organized

The document is organized as follows:

- Chapter 1, "Introduction to the VSE/ESA Turbo Dispatcher"

  This chapter provides an introduction to the VSE/ESA 2.1 Turbo Dispatcher.

- Chapter 2, "Planning for a Turbo Dispatcher Environment"

  This chapter describes the planning considerations for the VSE/ESA 2.1 Turbo Dispatcher.

- Chapter 3, "Operating a Turbo Dispatcher Environment"

  This chapter explains the operation aspects of a VSE/ESA 2.1 system with the Turbo Dispatcher.

- Chapter 4, "VSE/ESA with the Turbo Dispatcher under VM/ESA"

  This chapter provides the definitions of VSE/ESA 2.1 as a guest under VM/ESA in a multiprocessor environment.

- Chapter 5, "Description of the Test Environment"

  This chapter describes the system environment for the performance test runs described in the next chapter.

- Chapter 6, "Test Cases"

  This chapter documents the performance tests and results of a selected workload, using various system facilities in a VSE/ESA 2.1 multiprocessor environment.

- Chapter 7, "Guidelines for Using the VSE/ESA Turbo Dispatcher"

  This chapter provides guidelines for using the VSE/ESA 2.1 Turbo Dispatcher in a multiprocessor environment based on the results of the test runs described in the previous chapters.

- Appendix A, "IPL Load Parameter"

  This appendix shows the IPL load parameter and its defaults.

- Appendix B, "Setup Examples"

  This appendix gives setup examples for IPL and jobs.

Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this document.

- *Online Library, VSE Collection 09/95*, SK2T-0060-07
- *VSE/ESA 2.1 Turbo Dispatcher Guide and Reference*, SC33-6599-00
- *VSE/ESA 2.1 System Control Statements*, SC33-6613-00
- *VSE/VSAM 6.1 Commands*, SC33-6631-00
- *CICS/VSE 2.3 Resource Definition (Macro)*, SC33-0709-02
- *COBOL for VSE 1.1 Programming Guide*, SC26-8072-00
- *LE/VSE 1.1 Installation and Customization Guide*, SC26-8064-00
- *TPNS Planning and Installation*, SH20-2488
- *TPNS Script Generator Utilities*, SC30-3453
- *TPNS General Utilities*, SC30-3290

The following packages contain basic and useful information relating to this document and to VSE/ESA performance in general. The packages are available upon request from your IBM representative:

- *VE21PERF PACKAGE* - VSE/ESA 2.1 Performance Considerations

  A separate document about the performance aspects of the Turbo Dispatcher is included.
- *VE13PERF PACKAGE* - VSE/ESA 1.3/1.4 Performance Considerations
- *VSENEW11 PACKAGE* - VSE/ESA Newsletter, Third/Fourth Quarter 1995

The packages are stored on the IBM internal tools disk IBMVSE. IBM employees may obtain a copy of the packages by entering the following command from their local VM user IDs:

```
TOOLS SENDTO BOEVM3 VMTOOLS IBMVSE GET xxxxxxxx PACKAGE
```

where xxxxxxxx is the name of the package.

These packages are also available on Internet via the:

- FTP Server at ***ftp://lscftp.kgn.ibm.com/pub/vse/docs***
- VSE/ESA Home Page at ***http://www.ibm.de/go/d00000166***

The packages are in zipped format that can be down loaded, unzipped, and printed on a Postscript-capable printer:

- *VE21PERF.ZIP* - VSE/ESA 2.1 Performance Considerations
- *VE21TDP.ZIP* - Performance Considerations for the VSE/ESA 2.1 Turbo Dispatcher
- *VE13PERF.ZIP* - VSE/ESA 1.3/1.4 Performance Considerations
- *VSENEW11.ZIP* - VSE/ESA Newsletter, Third/Fourth Quarter 1995

The performance packages are updated periodically, normally when a new modification level of VSE/ESA is available.

International Technical Support Organization Publications

A complete list of International Technical Support Organization publications, known as redbooks, with a brief description of each, may be found in:

> *International Technical Support Organization Bibliography of Redbooks,* GG24-3070.

Especially the following redbook contains helpful information about the planning and usage of SQL/DS 3.5:

> *SQL/DS 3.5 Usage Guide*, SG24-4647

To get a catalog of ITSO redbooks, VNET users may type:

```
TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
```

A listing of all redbooks, sorted by category, may also be found on MKTTOOLS as ITSOCAT TXT. This package is updated monthly.

How to Order ITSO Redbooks

IBM employees in the USA may order ITSO books and CD-ROMs using PUBORDER. Customers in the USA may order by calling 1-800-879-2755 or by faxing 1-800-284-4721. Visa and MasterCard are accepted. Outside the USA, customers should contact their local IBM office. For guidance on ordering, send a PROFS note to BOOKSHOP at DKIBMVM1 or E-mail to bookshop@dk.ibm.com.

Customers may order hardcopy ITSO books individually or in customized sets, called BOFs, which relate to specific functions of interest. IBM employees and customers may also order ITSO books in online format on CD-ROM collections, which contain redbooks on a variety of products.

ITSO Redbooks on the World Wide Web (WWW)

Internet users may find information about redbooks on the ITSO World Wide Web home page. To access the ITSO Web pages, point your Web browser to the following URL:

```
http://www.redbooks.ibm.com/redbooks
```

IBM employees may access LIST3820s of redbooks as well. The internal Redbooks home page may be found at the following URL:

```
http://w3.itsc.pok.ibm.com/redbooks/redbooks.html
```

Subscribing to Internet Listserver

IBM redbook titles/abstracts are now available through Internet E-mail via the IBM Announcement Listserver. With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. All it takes is a few minutes to set up a profile, and you can get news (in ASCII format) from selected categories.

To initiate the service, send an E-mail note to:

announce@webster.ibmlink.ibm.com

with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

To obtain more details about this service, employees may type the following:

TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE

Note: INEWS users can select RelInfo from the action bar to execute this command automatically.

Acknowledgments

This project was designed and managed by:

Hanns-Joachim Uhl
VSE Development Laboratory, IBM Germany

The authors of this document are:

Osvaldo H. Di Carlo
IBM Argentina

Cheng Jing
IBM China

Johnny K. Yiu
IBM Canada, currently on assignment to IBM China

This publication is the result of a residency conducted at the International Technical Support Organization, Boeblingen Center.

Thanks to the following people for the invaluable advice and guidance provided in the production of this document:

Wolfgang Krämer
VSE Development Laboratory, IBM Germany

Ingolf Salm
VSE Development Laboratory, IBM Germany

Chapter 1. Introduction to the VSE/ESA Turbo Dispatcher

This chapter provides the reader with a short description of the concepts of the Turbo Dispatcher introduced in VSE/ESA 2.1.

This section contains information about:

- the general characteristics of the Turbo Dispatcher
- the concept of VSE/ESA multiprocessor support
- the memory layout used for the Turbo Dispatcher
- the partition balancing used by the Turbo Dispatcher

1.1 Turbo Dispatcher Overview

The dispatcher is that function of the VSE/ESA supervisor which dispatches units of work awaiting processing.

There are two types of dispatchers available with VSE/ESA 2.1:

- Standard Dispatcher
- Turbo Dispatcher

The **Standard Dispatcher** is the dispatcher also available with former VSE releases. It is still the default dispatcher in a VSE/ESA 2.1 system, and supports single processors, be they native, under VM or in a PR/SM LPAR.

VSE/ESA 2.1 offers for the first time an additional dispatcher called the **VSE/ESA Turbo Dispatcher** which is mainly designed to support multiprocessor environments.

The Turbo Dispatcher can utilize a multiprocessor by distributing the workload over several central processors (**CP**) within one Central Electronic Complex (**CEC**). The Turbo Dispatcher supports all the current IBM ESA/370 and IBM ESA/390 processors, such as the CMOS-based models S/390 9672 Parallel Enterprise Servers. Therefore, all these processors can now be used by VSE/ESA 2.1 also in native mode.

This manual uses the term multiprocessor for systems with more than one CP. The terms CP and processor are used interchangeably when referring to an individual processor of a multiprocessor.

In addition to the support of multiprocessors, the Turbo Dispatcher may also have advantages for uniprocessors. It provides better partition balancing between static and dynamic partitions compared to the Standard Dispatcher. This behavior will be explained later in more detail.

Further characteristics of the VSE/ESA Turbo Dispatcher are:

- VSE/ESA with the Turbo Dispatcher can run:
  - in native mode
  - in LPAR mode with PR/SM
  - under VM/ESA as a guest system

- The Turbo Dispatcher must be activated at IPL (Initial Program Load). Additional CPs can be started through the $0JCL procedure during system initialization or by the operator via the SYSDEF attention routine command.
- The support is transparent to existing user applications as well as to subsystems such as CICS/VSE and ACF/VTAM.

In general, the usage of the Turbo Dispatcher, either on a uniprocessor or a multiprocessor, should be taken into consideration because IBM will provide future dispatcher enhancements mainly for the Turbo Dispatcher.

1.2 VSE/ESA Multiprocessor Support

To ensure that a working VSE system is not influenced when running the Turbo Dispatcher, VSE has taken the following approach to implement the multiprocessor support.

The Turbo Dispatcher distributes work dynamically on a partition level, that is, it dispatches an entire partition to a CP waiting for work. Because of the dynamic distribution of work, no dedication of a CP to a specific partition can be done.

From a dispatching point of view, all CPs of a multiprocessor have equal rights. For example, all CPs can receive interrupts.

The program or application running in a partition consists, as seen by the Turbo Dispatcher, of many work units. One work unit is defined as a set of instructions processed from the point of selection by the Turbo Dispatcher until the next interrupt.

Only one work unit of a partition can be processed at a time. That is, no other work unit of the *same* partition can run on any other CP concurrently. The Turbo Dispatcher distinguishes between two types of work units:

1. **Parallel Work Units**

   Customer applications in a batch or online environment can be processed mostly as parallel work units.

   The parallel processing of a program is frequently interrupted for the processing of non-parallel work units, for example when system services are required.

2. **Non-Parallel Work Units**

   Non-parallel work units cannot be processed in parallel. Only one CP of a multiprocessor system can process a non-parallel work unit at any point in time. But at the same time other CPs can process parallel work units of other partitions. If there are multiple concurrent requests to execute non-parallel work units, these requests have to be queued.

   Typical examples of non-parallel work units are most system services and key-0 programs such as supervisor services, VSE/POWER or VTAM services.

Figure 1 shows a simplified example of how the Turbo Dispatcher processes a given workload consisting of a mixture of parallel and non-parallel work units. The example assumes a multiprocessor with two CPs (CP 0 and CP 1) and three jobs (A, B and C) running in different partitions.

Job A has the highest, job B the second highest and job C the lowest priority. Each job consists of three consecutive work units: A1, A2 and A3, for example. Both types of work units are to be processed: parallel (**P**) and non-parallel (**N**). In the example, it is assumed that each work unit consumes the same amount of CPU time and no interrupts are to be processed. It demonstrates

- how the Turbo Dispatcher distributes work units on partition level, and
- how the occurrence of non-parallel work units influences the job sequence.

	UNIPROCESSOR		MULTIPROCESSOR		
				CP 0	CP 1
JOB A	A1 (N)	Step 1	Step 1	A1 (N)	B1 (P)
	A2 (N)	Step 2	Step 2	A2 (N)	
	A3 (P)	Step 3	Step 3	B2 (N)	A3 (P)
JOB B	B1 (P)	Step 4	Step 4	B3 (P)	C1 (N)
	B2 (N)	Step 5	Step 5	C2 (P)	
	B3 (P)	Step 6	Step 6		C3 (P)
JOB C	C1 (N)	Step 7			
	C2 (P)	Step 8			
	C3 (P)	Step 9			

Figure 1. Processing Steps for Jobs A, B and C

On a uniprocessor either with the Standard or the Turbo Dispatcher all job steps have to be serialized. A total of nine steps is needed.

On a multiprocessor under the control of the Turbo Dispatcher, the three jobs are to be processed in six steps only:

1. Since jobs A and B have a higher priority than job C, CP 0 starts processing with non-parallel work unit A1 and CP 1 starts with parallel work unit B1.
2. The next non-parallel work unit A2 is selected on CP 0; work units B2 and C1 are available but both are non-parallel and cannot be selected because only *one* non-parallel work unit can be active at a time.
3. Non-parallel work unit B2 and parallel work unit A3 are selected by CP 0 and CP 1; non-parallel work unit C1 still has to wait due to its lower priority.
4. Job A has completed. Parallel work unit B3 and non-parallel work unit C1 are selected now by CP 0 and CP 1.
5. Job B has also completed. CP 0 processes the parallel work unit C2. Parallel work unit C3 is to be queued because only one work unit of a partition can be executed at any one time.
6. CP 1 processes the last parallel work unit C3.

This example shows that whenever a non-parallel work unit is executing on one CP, all other concurrent non-parallel work units have to be queued. When this happens, a CP can be selected to execute a parallel work unit of another partition. After the CP has finished the non-parallel work unit, the dispatcher decides, depending on the priority, which of the available work units is to be processed next.

1.3 Memory Layout of a VSE/ESA Multiprocessor Environment

In a VSE/ESA multiprocessor environment, all CPs share the real storage and have access to the private and shared areas of VSE/ESA 2.1 including the supervisor. Basically, there is no change in the VSE/ESA 2.1 system and address space layout.

Because information in the first 4KB page of the storage (low core) is specific to a CP, for each CP a 4KB page, called **Prefix Page**, was introduced with the Turbo Dispatcher. This means that the first 4KB page is unique to each CP.

Figure 2 on page 5 shows the usable storage of two CPs, where CP 0 is executing a CICS/VSE partition and CP 1 a VTAM partition. This partition assignment is to be seen as a snapshot only since it may change from the processing of one work unit to the next.

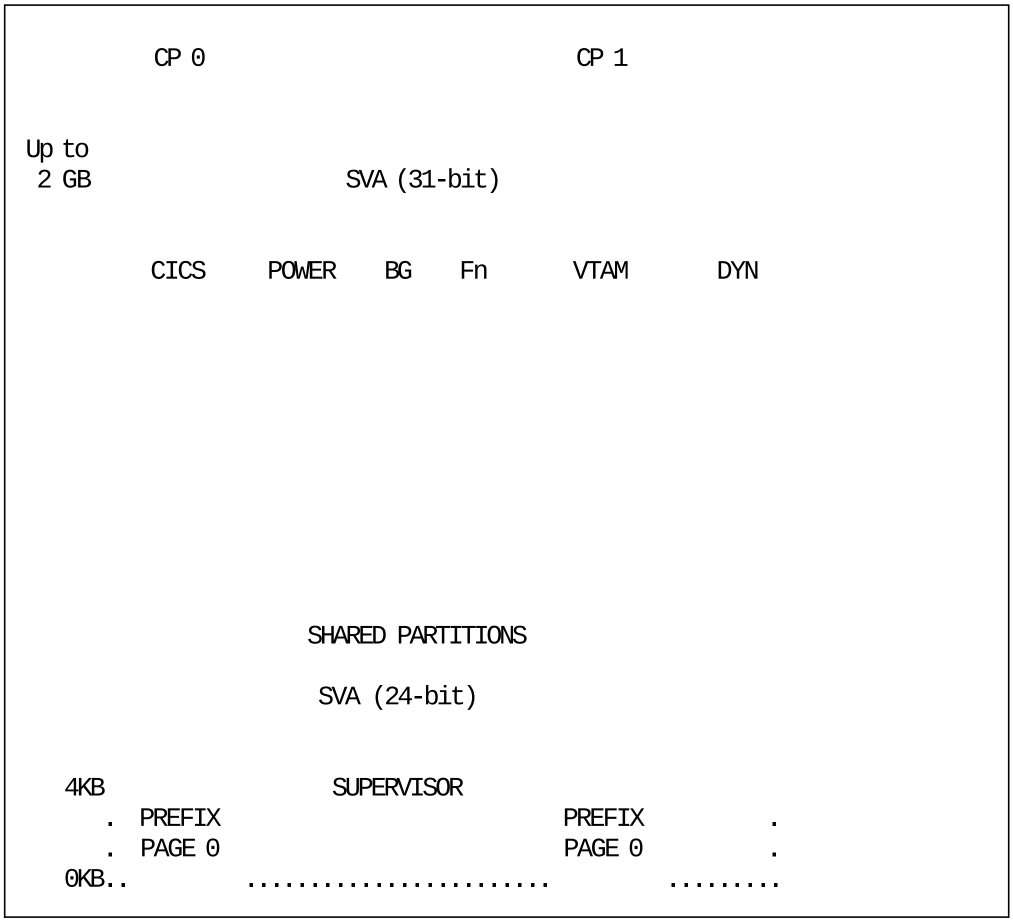

Figure 2. Sharing of VSE/ESA Virtual Storage in a Multiprocessor System

If during a non-parallel work unit control blocks in the prefix page of the current CP were changed, these changes will become effective on all other CPs after the non-parallel work unit is finished.

1.4 Partition Balancing

The Turbo Dispatcher provides an improved partition balancing algorithm, which gives each partition, be it static or dynamic, equal weight in the balanced group.

If the Standard Dispatcher is active, each dynamic partition class receives the same time slice as each static partition.

In the Turbo Dispatcher environment with the new partition balancing algorithm, the time slice for any dynamic partition is equal to that of a static partition.

An example shows how the changed weights can influence the VSE system. The example requests partition balancing for dynamic partitions of classes C and D, and the static partitions BG, F5 and F7:

```
PRTY C=D=BG=F5=F7
```

It is assumed that currently five dynamic partitions of class C are active and 10 partitions of class D.

The following summary shows how the sizes of the time slices change depending on whether the Standard Dispatcher or the Turbo Dispatcher is used.

STANDARD DISPATCHER

C	D	BG	F5	F7	
1/5	1/5	1/5	1/5	1/5	each class/static partition
1/25	1/50	1/5	1/5	1/5	each individual partition

TURBO DISPATCHER

D	C	BG	F5	F7	
1/18	1/18	1/18	1/18	1/18	each individual partition

Figure 3. Dispatcher Time Slices

In the Standard Dispatcher environment any partition of class D will get 1/50 of the time slice, in the Turbo Dispatcher environment 1/18 of the time slice.

The example shows that the Turbo Dispatcher with the new partition balancing algorithm provides more CPU time to a dynamic partition in a partition balancing group than with the Standard Dispatcher.

The time slice can be dynamically changed by the MSECS command.

For further details about the MSECS and PRTY commands see the *VSE/ESA System Control Statements* manual.

Chapter 2. Planning for a Turbo Dispatcher Environment

It is the type and setup of a workload of a VSE/ESA 2.1 customer environment that determines whether a specific multiprocessor can be exploited with the Turbo Dispatcher.

The objective of this chapter is to provide information how to determine, evaluate and plan an environment suitable for use by the Turbo Dispatcher.

This section discusses in detail:

- The non-parallel share
- The maximum number of exploitable CPs
- The planning for partitions
- The reduction of the non-parallel share

2.1 Non-Parallel Share and Number of CPs

Efficient use of the Turbo Dispatcher depends mainly on the number and type of work units that can be dispatched for processing. When more parallel work units are available for concurrent dispatching, more CPs can be utilized in parallel.

The Turbo Dispatcher identifies work units and their status and initiates their processing according to the following rules:

- If a CP processes a work unit of a particular partition, no other CP can execute any other work unit of that partition.
- For partitions that attach VSE subtasks, no other tasks of the same partition can be processed on any other CP when one task of that partition is already active.
- Only one CP can process a non-parallel work unit at any one time. This means that during this time no other CP can process a non-parallel work unit even if it belongs to another partition. For example programs running with key-0 are executed non-parallel.

To find out the maximum number of CPs that can be exploited, first the non-parallel share of a given workload has to be determined.

The **Non-Parallel Share** is the ratio of non-parallel work units to all work units in the system. In other words the non-parallel share is the percentage of the total CPU time used for running non-parallel work units.

With VSE/ESA 2.1, the QUERY TD command can be used to find the non-parallel share of a given workload. For workstation users, *VSE Workdesk* offers the *CPU Activity* selection panel. Both functions can even be used in a single CP environment but only if the Turbo Dispatcher is active. For further details about the QUERY TD command see section 3.2.1, "QUERY TD" on page 13.

After determining the non-parallel share of a given workload the maximum number of exploitable processors shown in Table 1 on page 8 can be used for planning a Turbo Dispatcher environment. The meaning of the values in the table is as follows:

1. Non-parallel share

   As discussed above these values show the ratio of non-parallel CPU time to total CPU time provided by the QUERY TD command. The CPU time is the time spent by all CPs of a multiprocessor for a given workload.

2. Maximum number of exploitable CPs

   This value is calculated from the value for the non-parallel share of the workload as follows:

   number of CPs = 0.9/non-parallel share

   The value 0.9 is used here to take into account the delay caused by waiting for the non-parallel state. The 0.9 indicates as well that, in total, up to 90% of the time the NP state can be used to execute non-parallel work units.

Non-Parallel Share	0.20	0.25	0.30	0.35	0.40	0.45	0.50	0.55
Maximum number of exploitable CPs	4.5	3.6	3.0	2.6	2.2	2.0	1.8	1.6
Note: • The table is extracted from the ′VSE/ESA 2.1 Turbo Dispatcher Guide and Reference - SC33-6599′								

Table 1. Relationship of Non-Parallel Share and Exploitable CPs

With the non-parallel share for a given workload further planning considerations can be made to see how to exploit the Turbo Dispatcher more efficiently. These considerations are described now in more detail.

2.2 Planning for Partitions

Since the Turbo Dispatcher distributes work units on a partition level, any single partition can exploit at most the processing power of a single CP only, even if other CPs are waiting for work.

A single partition normally exploits less than the processing power of a single CP, if other partitions frequently cause that partition to wait for the non-parallel state. In practice however, the partition with the highest workload dictates the actual processing requirement for a single CP, and in many cases exceeds the capabilities of such a processor.

Especially large CICS/VSE applications, running in a single partition, can create a higher workload than can be processed by a single CP in a multiprocessor environment. Whenever possible, such applications should be split and distributed over more than one partition, thus enabling the Turbo Dispatcher to distribute the workload over more than one CP for processing.

How a split of a large CICS/VSE partition in a multiprocessor environment could be done, is shown in the following two examples:

1. Using independent CICS/VSE partitions

   If possible, the splitting of one large CICS/VSE partition into multiple independent CICS/VSE partitions is the easiest way to distribute the resulting workload.

2. Using CICS/VSE MRO

   The CICS/VSE Multiregion Operation (MRO) can be used together with Transaction Routing (TR) and Function Shipping (FS) to split one large CICS/VSE partition into multiple smaller CICS/VSE partitions. These special-purpose CICS/VSE partitions can be AOR (Application Owning Region), TOR (Terminal Owning Region), and FOR (File Owning Region), or combinations of these. This can be achieved by:

   - Using TR to several target CICS/VSE partitions

     When using TR, transactions can be routed to more than one target CICS/VSE partition. AOR and FOR can be combined and run in a single partition.

   - Using FS to a target FOR

     When using FS, file requests to a file owning region (FOR) may be started from more than one CICS/VSE partition. In that case TOR and AOR are combined in one or more partitions while FOR processing is done in a separate partition.

In any case the MRO overhead for TR and FS has to be taken into account. For further details see the specific performance documentation.

Another possibility to split the workload across different partitions is to use the Dynamic Database Switching facility of SQL/DS 3.5. It allows a CICS transaction to dynamically switch between different SQL/DS application servers with different Logical Units of Work.

The VSE/ESA 2.1 user may take advantage of this facility to distribute the load of a single SQL/DS application server across multiple SQL/DS application servers by using more partitions in the system. These would then enable the Turbo Dispatcher to distribute the workload more effectively to the available CPs.

Therefore, distributing the workload to multiple SQL/DS application servers in a VSE/ESA 2.1 environment provides another tuning potential on multiprocessor systems.

In general the following should be kept in mind when planning a system environment for the Turbo Dispatcher:

- More partitions in the system mean more work units are available for distribution among the available CPs.
- The lower the share of non-parallel work, the smaller are the delays caused by waiting for the non-parallel state and the higher the number of CPs that can be exploited, provided enough partitions are available.

How the non-parallel share can be reduced is discussed below.

2.3 Planning for Applications

The majority of existing application programs can run unchanged with the Turbo Dispatcher. Only those applications that interface with internal supervisor data may need modification to reflect the changes in the supervisor needed for the Turbo Dispatcher. This mainly applies to system management products such as performance management tools.

To most current level of the used products and the Turbo Dispatcher itself should be installed to ensure appropriate results. User programs written in a high level language such as COBOL or PL/I, usually do not need any modification.

2.4 Reducing Non-Parallel Work Units

As described in the previous sections, the non-parallel share is the most important factor to determine whether a current workload can exploit a multiprocessor. This assumes that enough partitions are active at a time. To make the workload more suitable for a multiprocessor, the reduction of non-parallel work units is important.

The share of non-parallel work can be reduced mainly by decreasing the number of I/Os. This can be achieved, for example, by making use of 31-bit addressing which allows to have more Data In Memory (DIM).

Examples for Data In Memory are:

- Using more and larger VSAM buffers above the 16MB line.
- Using more advanced programming techniques such as data space usage and access register mode.
- Using S/390 Software-/Hardware-Assisted Data Compression.

More Data In Memory means less I/O operations to external storage devices and thus less system code that must be processed as non-parallel work.

The share of non-parallel work depends on the types of applications that make up the workload. As a general rule, I/O intensive applications may have a non-parallel share of up to 50% whereas less I/O intensive online applications may have a non-parallel share of about 25% to 40%. But not only this I/O intensity determines the non-parallel share; the following should also be taken into consideration when planning to reduce the non-parallel share of a workload:

- Virtual Disks

  The support for virtual disks runs mostly non-parallel. Heavy use of virtual disks may thus increase the non-parallel share accordingly.

- NPA Operand

  Using the NPA operand in the EXEC statement will cause the program to run as non-parallel. For further details on the NPA operand see 3.5, “NPA” on page 18.

With the QUERY TD command it is possible to check the change in the non-parallel share of a workload when the system environment is altered.

If migration planning has to be done based on VSE/ESA V1, the non-parallel share cannot be derived. A high I/O activity usually indicates that the share of non-parallel work units will be high which would reduce the CP exploitation capability of the Turbo Dispatcher for a multiprocessor accordingly.

To ensure that there are no unrealistic expectations, the non-parallel share should be verified as described above after the migration to VSE/ESA V2 is completed.

Chapter 3. Operating a Turbo Dispatcher Environment

The objective of this section is to provide the necessary information to control a VSE/ESA 2.1. Turbo Dispatcher environment.

The following is described in detail:

- IPL - how to activate the Turbo Dispatcher
- Commands - to control the Turbo Dispatcher environment
- Display System Activity - to check the system activity
- SDAID - restrictions when using SDAID
- NPA - to force a program to work non-parallel

3.1 IPL

The operator must activate the Turbo Dispatcher with the IPL load parameter by specifying a **T** in the fifth position. The load parameter can be entered as follows:

- For VSE/ESA 2.1 **native**

  In the load parameter field of the program load panel.

- For VSE/ESA 2.1 as a **guest system** under VM

  ```
  IPL cuu LOADPARM  ....T
  ```

The detailed meaning of all LOADPARM parameters is shown in Appendix A, "IPL Load Parameter" on page 47.

During IPL only one CP is active. After IPL, additional CPs can be activated with the SYSDEF TD command either as an attention routine command or included in the $0JCL procedure.

As shown in Figure 4 on page 12, it is possible to switch from the VSE/ESA 2.1 Turbo Dispatcher to the Standard Dispatcher and vice versa by IPLing anew.

Note: The Standard Dispatcher is the default in VSE/ESA 2.1.

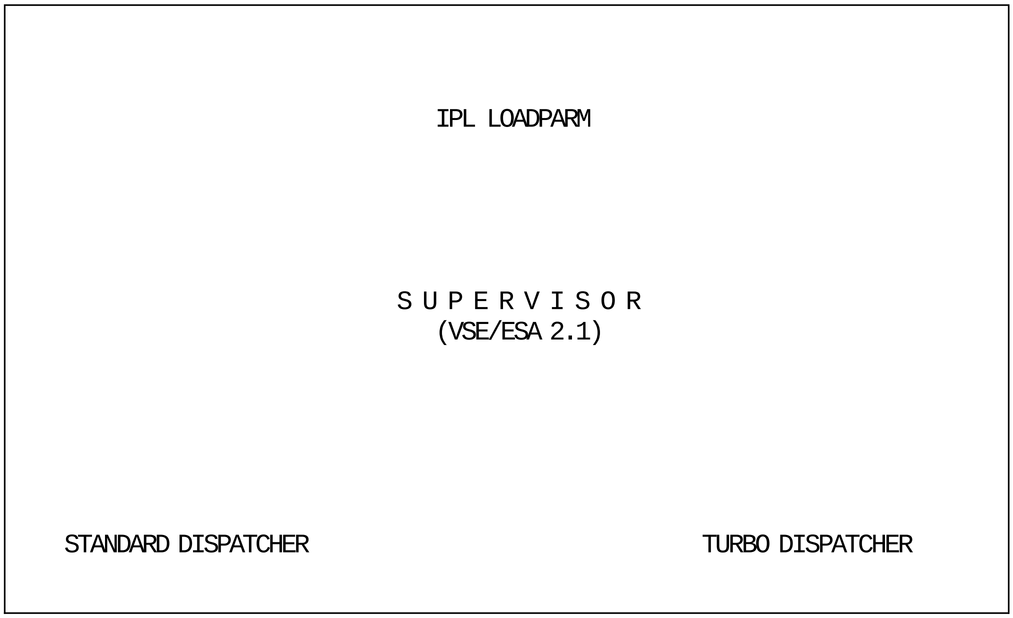

Figure 4. Dispatcher Selection

The following message is issued during IPL to inform the user that the Turbo Dispatcher is activated:

0J65I TURBO DISPATCHER ACTIVATED

The Turbo Dispatcher requires VSE Job Accounting (JA). The default setting for this parameter in the IPL SYS command is JA=NO.

If JA=....is not specified in the IPL SYS command and the Turbo Dispatcher is active, the system sets the value to:

SYS JA=YES

Message 0J64I is issued to inform the user of the changed setting:

0J64I SYS JA=YES ASSUMED BECAUSE OF TURBO DISPATCHER ACTIVATION

A specification of JA=NO is not accepted. If specified, the setting is ignored and the message 0J64I is issued.

3.2 Commands

To control a Turbo Dispatcher environment, VSE/ESA 2.1 has introduced new parameters in the QUERY and SYSDEF commands or statements. Because of the introduction of the prefix page, specific restrictions now apply to the commands DSPLY and ALTER.

Both the new commands and the restrictions are described in more detail below.

3.2.1 QUERY TD

To query the status of VSE/ESA on a multiprocessor the following attention routine command was introduced with VSE/ESA 2.1:

```
QUERY TD
```

Figure 5 shows the result when the QUERY TD command was issued on a six-way processor where only three processors were started:

```
QUERY TD


AR 0015  CPU   SPIN_TIME     NP_TIME  TOTAL_TIME NP/TOT
AR 0015   04        2603      306410      562970  0.544
AR 0015   00    INACTIVE
AR 0015   01        2834      220094      484201  0.454
AR 0015   02    INACTIVE
AR 0015   03    INACTIVE
AR 0015   05        2009       66565      176869  0.376
AR 0015       ----------------------------------------
AR 0015 TOTAL       7446      593069     1224040  0.485
AR 0015
AR 0015  ELAPSED TIME SINCE LAST RESET:      60516670
AR 0015 1I40I  READY
```

Figure 5. Example of Query TD Command Output

The information displayed by the QUERY TD command has the following meaning:

CPU
The CP address. It is defined by the available CPs seen by VSE.

The first line shows the CP number from where IPL was performed, the next lines show those CPs which can be activated or stopped by the SYSDEF TD command.

SPIN_TIME
The time in milliseconds during which the CP was within an instruction loop waiting for a resource occupied by another task.

If a CP has not yet started or has been stopped, the character string INACTIVE is displayed for that CP.

NP_TIME
The time in milliseconds during which the CP processed non-parallel work units.

TOTAL_TIME
The time in milliseconds during which the CP has processed either parallel or non-parallel work units. This means that the TOTAL_TIME value includes the NP_TIME value.

Note: The TOTAL_TIME does not include the SPIN_TIME value.

NP/TOT The non-parallel share. This is the ratio of non-parallel time to total time; that is, the quotient of NP_TIME and TOTAL_TIME.

For more details about the non-parallel share see Chapter 2, "Planning for a Turbo Dispatcher Environment" on page 7.

TOTAL Shows the total sum for each column. The total NP/TOT value is the quotient of the total values for NP_TIME and TOTAL_TIME.

ELAPSED TIME SINCE LAST RESET
Shows in milliseconds the time passed since the last reset of CP related information. Such a reset occurs whenever a SYSDEF TD command or statement is being processed.

Note that all values apply to the time since the last reset shown in ELAPSED TIME SINCE LAST RESET.

3.2.2 SYSDEF TD

With the new SYSDEF TD command or statement, CPs can be started and stopped. Turbo Dispatcher information can also be reset.

3.2.2.1 Starting CPs

There are two ways to start CPs:

1. In the startup procedure $0JCL the following statements can be included:

   // SYSDEF TD,START=ALL to start all available CPs

   // SYSDEF TD,START=*cpuaddr* to start a specific CP

2. At the VSE/ESA 2.1 console the following attention routine (AR) commands can be entered:

   SYSDEF TD,START=ALL to start all available CPs

   SYSDEF TD,START=*cpuaddr* to start a specific CP

ALL means that all CPs of a multiprocessor that are not yet active are to be activated. *cpuaddr* means that a single CP identified by *cpuaddr* has to be activated. *cpuaddr* can be any hexadecimal value from X′00′ to X′FF′ and possible candidates are listed by the QUERY TD command.

A START request implies a RESETCNT request.

The only way to use the SYSDEF TD,START statement is in the $0JCL procedure. If the statement is issued in any job or procedure running in any static or dynamic partition (except BG), the following message is displayed:

1UV7D SYSDEF ONLY ALLOWED IN BG OR AS ATTENTION COMMAND

If the SYSDEF TD,START statement is issued in the BG partition the following message is displayed:

1YK3D ASI NOT ACTIVE: START PARAMETER NOT ALLOWED

3.2.2.2 Stopping CPs

All CPs or a single CP can be stopped. The following attention routine (AR) commands are available:

SYSDEF TD,STOP=ALL   to stop all CPs

SYSDEF TD,STOP=*cpuaddr*   to stop a specific CP

ALL means that all CPs of a multiprocessor that are not yet stopped are to be stopped. *cpuaddr* means that a single CP identified by *cpuaddr* has to be stopped. *cpuaddr* can be any hexadecimal value from X′00′ to X′FF′ and possible candidates can be derived from the QUERY TD command. The IPL CP cannot be stopped.

A STOP request implies a RESETCNT request.

The SYSDEF TD,STOP command is only allowed as an AR command. If it is issued in any job or procedure, the following message is displayed:

1YA4D NOT ALLOWED IN AR OR JCL

If the SYSDEF TD,STOP command is issued for the IPL-CP the following message is displayed:

1YH6I CPU(s) COULD NOT BE STOPPED RC=08 REASON=10

RC=08, RS=10 means : It is not allowed to stop a CP from which the system has been IPLed.

CPs stopped with the SYSDEF TD command will get the following special disabled wait PSW:

PSW : 000A0000 00001111

Under a VM system the following message is displayed when stopping a CP with the SYSDEF TD,STOP command:

HCPGIR450W CP ENTERED; DISABLED WAIT PSW 000A0000 00001111

3.2.2.3 Resetting Turbo Dispatcher Information

Turbo Dispatcher information can be reset with the following attention routine (AR) command:

SYSDEF TD,RESETCNT

This command is executed automatically when a CP is stopped or started with the SYSDEF TD,START or STOP command.

The reset should be used, for example, before issuing a QUERY TD command to get a defined starting point to evaluate the non-parallel share of a workload.

3.2.2.4 Using the SYSDEF Command in a VM/ESA Environment

If the VSE console is the CP console (not dialed), it is recommended to first issue the CP command **SET EMSG ON** before issuing the SYSDEF command. This is to reduce heavy message traffic that may occur.

3.2.3 DSPLY

The **DSPLY** command allows the operator to display 16 bytes of virtual storage, starting at a specified hexadecimal address, on the device assigned to SYSLOG.

If the VSE/ESA 2.1 Turbo Dispatcher is active and at least two CPs are started, the address range from X′0′ to X′FFF′, the prefix page, is no longer unique. Hence, if the hexadecimal address specified with the DSPLY command is less than X′1000′, the output of the DSPLY command is random and shows storage specific for one active CP.

To avoid the wrong output, additional CPs should be stopped before issuing the DSPLY command for addresses below X′1000′.

For further details about the DSPLY command see the *VSE/ESA System Control Statements* manual.

3.2.4 ALTER

The **ALTER** command allows the operator to alter 1 to 16 bytes of virtual storage, starting at a specified hexadecimal address.

If the VSE/ESA 2.1 Turbo Dispatcher is active and at least two CPs are started, the address range from X′0′ to X′FFF′, the prefix page, is no longer unique. Hence, if the hexadecimal address specified with the ALTER command is less than X′1000′, the storage area to be altered could be selected randomly depending on which CP is active. It also involves the risk that system-relevant data is destroyed. Because of this the system will reject an ALTER command on the prefix page if more than one CP is active.

The following message is displayed to inform the user that the ALTER command is not possible on the prefix page:

1I37I UPDATE ON PREFIX PAGE NOT POSSIBLE IN MP ENVIRONMENT

In this message, MP stands for multiprocessor.

However, if necessary, the ALTER command can be used in uniprocessor mode by stopping all additional CPs first. After returning to multiprocessor mode, the changed information is automatically transferred to the other started CPs.

For further details about the ALTER command see the *VSE/ESA System Control Statements* manual.

3.3 Display System Activity

An example of the Interactive Interface Display System Activity dialog (fastpath '361') for CPU usage in a VSE/ESA multiprocessor environment is shown in Figure 6. The system is running two production CICS and one supervisor generation job.

```
IESADMDA            DISPLAY SYSTEM ACTIVITY           224 SECONDS  11:19:54
*--- SYSTEM (CPUS ACTIVE:  3 ) ---* *------------ CICS : DBDCCICS ------------
|CPU      : 167%   I/O/SEC:  392  | |NO. TASKS:   5,455   PER SEC    :   0.1
|PAGES IN :    0   PER SEC:    *  | |ACTIVE TASKS:    6   SUSPENDED  :     3
|PAGES OUT:    0   PER SEC:    *  | |MOST ACTIVE :    6   PAGES AVAIL:   509
*--------------------------------* *----------------------------------------
PRIORITY: Z,Y,K,J,H,G,P,F6,FB,FA,F7,BG,F8,F9,F4=F5,F2,F3,F1

 ID S JOB NAME  PHASE NAME  ELAPSED    CPU TIME   OVERHEAD  %CPU       I/O
 F1 1 POWSTART   IPWPOWER   31:01:05    394.62     168.68     8%   108,984
 F3 3 VTAMDSW    ISTINCVT   31:00:49    311.15      81.89        1,314,723
 F2 2 CICSICCF   DFHSIP     31:00:27    361.25     344.49          119,220
 F5 4 PRODCICS   DFHSIP     00:31:29    782.24     122.71    65%   211,827
 F4 5 PRODCIC2   DFHSIP     00:31:23    776.89     121.34    64%   213,213
 F9 6 SUPRGEN    ASMA90     00:09:27    130.14      26.67    30%     8,684
 F8 7 <=WAITING FOR WORK=>                 .11        .06               18
 BG 8 <=WAITING FOR WORK=>                 .00        .00                2
 F7 9 <=WAITING FOR WORK=>                 .02        .01               18
 FA A <=WAITING FOR WORK=>                 .02        .01               18
 FB B <=WAITING FOR WORK=>                 .02        .01               18
 F6 0 <=WAITING FOR WORK=>                 .02        .01               18
PF1=HELP                    3=END       4=RETURN     5=DYN.PART   6=CPU
```

Figure 6. Example of the Display System Activity Screen

The second line in Figure 6 shows the number of CPs activated, while the third line displays the total CP usage (167%) which is the sum of the individual partition utilization and therefore, the total CPU activity of all active CPs. Naturally for a single CP the maximum value is 100%.

In this figure the interval timer is 224 seconds. For this panel the interval time was changed from 15 to 0 seconds to display a screen update only when the ENTER key was pressed. For further details of how to change the interval time, refer to *VSE/ESA Administration* (SC33-6505).

3.4 SDAID

SDAID can only be activated in single CP mode. Therefore, as long as SDAID is active, using multiple processors is not possible.

The **STRTSD** command will be rejected if more than one CP is active. In that case the following messages will be issued:

4C49I MULTIPROCESSING IS ACTIVE

4C05I PROCESSING OF STRTSD FAILED

The operator has to use the SYSDEF TD,STOP=ALL command to stop all additional CPs.

When SDAID is active, any attempt to start a CP will be rejected with the following message:

1YH5I CPU(s) COULD NOT BE STARTED RC=08 REASON=0F

RC=08, RS=0F means: The START request is rejected because SDAID is currently active.

3.5 NPA

The parameter **NPA** (Non-Parallel Application) is introduced as a new operand of the EXEC command or statement.

The NPA parameter defines a job step to be executed non-parallel. This operand can become necessary, when the VSE/ESA 2.1 Turbo Dispatcher is active and at least one additional CP has been started.

Figure 7 shows the new flow diagram of the EXEC command or statement.

```
      EXEC
//                 progname    ,REAL   ,SIZE=  nK
            PGM=                               mM
                                               AUTO
                                               (AUTO,   nK)
                                                        mM
                                               phasename
                                              (phasename,   nK   )
                                                            mM


 ,GO        ,PARM=¢value¢      ,DSPACE=  nK
                                         mM
                               ,TRACE
                               ,NPA


          EXEC REXX=procname

  //                          ,SIZE=    nK
                                        mK
                                        phasename
                                        (phasename,   nK    )
                                                      mK

  ,PARM=¢value¢           ,NPA
```

Figure 7. EXEC Format

An application program usually consists of parallel and non-parallel work units. It is the Turbo Dispatcher's task to determine the category to which a work unit belongs and to handle it accordingly. In particular problem situations it may be desirable to interfere and enforce non-parallel processing for a program. For this purpose the NPA operand is available.

The NPA operand should be used for problem solving only and may be useful in cases such as the following:

- A program runs correctly on a single CP but not on a multiprocessor.
- Two programs run and communicate correctly on a single CP but not on a multiprocessor. This may be a synchronization problem.

NPA can also be specified in the VSE/ICCF job entry statement:

/LOAD *programname* NPA

The following diagram shows the complete flow diagram of the /LOAD job entry statement in VSE/ICCF.

```
/LOAD phasename
                              PARM=¢parameter¢        NPA
```

Figure 8. /LOAD Job Entry statement (VSE/ICCF)

Chapter 4. VSE/ESA with the Turbo Dispatcher under VM/ESA

VSE/ESA 2.1. with the Turbo Dispatcher can run as a guest under VM/ESA 1.2.0 and all subsequent releases. Note that VSE/ESA 2.1 runs with supervisor mode ESA only.

VM/ESA provides multiprocessing support as follows:

- Real Multiprocessing

  An individual real processor (CP) can be dedicated to a guest system using the DEDICATE option of the CPU statement in the directory entry for the guest machine.

  This type of multiprocessing is not possible on a uniprocessor.

- Virtual Multiprocessing

  Guest systems can have multiple virtual processors assigned using the CP DEFINE command.

  This type of multiprocessing is possible even on a uniprocessor.

VM/ESA distinguishes between the following types of processors:

- Master and Alternate processors

  These are both real processors (CPs). The master processor is required by VM/ESA for certain VM CP work. It is usually the CP from which IPL was performed. The master processor cannot be dedicated to any guest system. Any other real processor is called an alternate processor.

- Dedicated processor

  This is a real processor that is dedicated to a single guest system for exclusive use.

- Base processor

  This is a virtual processor required by VM/ESA to manage the resources of a guest system.

Figure 9. VM/ESA Using a Multiprocessor with two CPs

An example will be used to demonstrate the steps required to run VSE/ESA 2.1 with the Turbo Dispatcher under VM/ESA. It is assumed that the hardware system has two real CPs. The following points are also assumed for this example:

1. The VSE/ESA Turbo Dispatcher guest is to be defined as a V=R Guest (or V=F) to utilize the advantages of a preferred guest system.
2. Two CPs are to be defined to VSE/ESA, because:
   - One CP is not sufficient for the guest's workload.
   - Defining more than two CPs results in poorer guest performance.
   - No dedication of one of the two real CPs is to be done because no more than two real CPs are available. If there are enough CPs available, the DEDICATE option could be used to dedicate a specific CP to a guest VSE/ESA.

Given the above, two examples of tailoring the VM/ESA directory entry for this guest show how to define the number of virtual processors.

1. Specify the maximum number of CPs in the MACHINE statement of the directory entry and use the VM CP DEFINE command to define the individual CPs.

   In the directory entry:

   ```
   MACHINE ESA 2
   ```

   Under VM CP:

   ```
   DEFINE CPU 00
   DEFINE CPU 01
   ```

   These commands define the virtual processors with virtual addresses 00 and 01. In VM/ESA virtual processors with addresses from 00 to 64 can be defined.
2. Specify the maximum number of CPs in the MACHINE statement and include CPU statements in the directory entry of the guest machine to allow the definition of individual characteristics of a virtual CP.

   In the directory entry:

   ```
   MACHINE ESA 2
   CPU 00 CPUID xxxxxx NODEDICATE
   CPU 01 CPUID xxxxxx NODEDICATE
   ```

   where xxxxxx is a six-digit hexadecimal CPU identification number.

Note

- The definition of more virtual processors than available real processors to a guest, will result in poor performance. Therefore, there are no performance reasons to define more than one virtual processor to a guest system on a uniprocessor, apart from, for example, for functional test reasons.
- To gain optimal performance, only as many processors (CPs) as really required for processing the total workload of a single VSE/ESA Turbo Dispatcher guest system should be defined to the VSE/ESA guest system.

Chapter 5. Description of the Test Environment

This chapter describes the test environment for the test runs conducted in Chapter 6, "Test Cases" on page 29. It contains the descriptions of the following components:

- System Environment
- TPNS (Teleprocessing Network Simulator)
- Application and Data Environment
- VSE/ESA System Setup

5.1 System Environment

All of the test runs were performed in a native VSE/ESA 2.1 environment, since it allows a controlled setup to test the change of one system parameter at a time with all the performance implications.

The VSE/ESA 2.1 system under test ran on an IBM 9672-R61 processor (six CPs) with 256MB of central storage. The test data resided on a total of 28 read-cached IBM 9345 volumes. The DASD units were distributed evenly over two IBM 9343-CC4 DASD controllers with 64MB of controller cache each. Each IBM 9343 DASD controller was connected to the VSE/ESA 2.1 system via three parallel channels.

The network and workload simulator, TPNS (Teleprocessing Network Simulator), ran as a VM guest machine on an IBM 9221-421 with 512MB of main memory, the Driver System. It was connected to the VSE/ESA 2.1 system via an IBM 3745 Channel Adapter through a VTAM-VTAM connection.

The workload and terminals simulated by TPNS were seen by the VSE/ESA 2.1 system as remote Logical Units (LUs) owned by the VM VTAM attached via a Channel-to-Channel (CTC) connection.

Figure 10 on page 24 outlines the test environment.

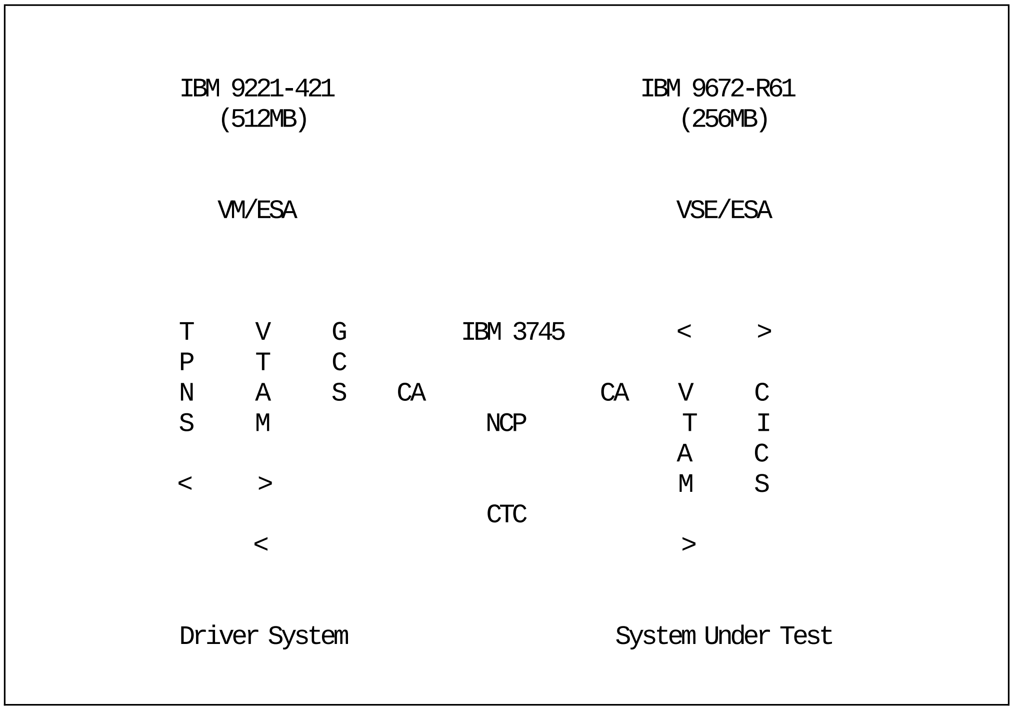

Figure 10. The Test Environment

5.2 TPNS - Network and Workload Simulator

The terminals and network traffic in the test runs were simulated by the IBM program product, Teleprocessing Network Simulator (**TPNS**), program number 5688-121. This program product runs under VM or MVS as a VTAM or IBM TCP/IP application. The key functions of TPNS are:

- Simulation of the network configuration and data traffic
- Definition of terminals and transactions to the target subsystem. The target subsystem may be CICS, IMS, TSO or DB2
- Definition of terminal operators′ think time
- Provision of the transaction throughput and response time statistics

5.3 Application and Data Environment

RAMP-C (Requirements Approach to Measuring Performance in COBOL) is a workload designed to work as a typical online application, but it is not patterned after a specific business application.

The test runs performed used the RAMP-C workload with VSAM, as the majority of VSE/ESA installations have VSAM as their preferred file access method.

The RAMP-C workload uses four classes of CICS applications and five sets of VSAM files.

The four classes of CICS applications differ in their degree of complexity and file accesses:

- Class 1 application issues two logical file I/Os

- Class 2 application issues seven logical file I/Os
- Class 3 application issues 23 logical file I/Os and calls subroutines
- Class 4 application issues 41 logical file I/Os and calls subroutines

The five sets of VSAM files are identified as A, B, C, D and E files:

- C files are VSAM ESDS files to be used as application log files. Records are appended to the end of the files by the transactions, and the file size grows with the volume of transactions. Therefore, C files are refreshed before each test run.
- The other sets of files are VSAM KSDS files used as data files. Each file has 30,000 records after it is initially loaded, and has an approximate size of 7.5MB.
- A files have a logical record length of 209 bytes. C files have a logical record length of 15 bytes. The other files have a logical record length of 244 bytes.
- All the files are distributed across six IBM 9345 volumes per CICS, to spread the DASD I/Os over different actuators.

A total of 84 files per CICS was used in the test runs. This includes four C files and 20 sets of A, B, D and E files. Table 2 summarizes the relationships between the files and the applications:

File	Appl. Class 1	Appl. Class 2	Appl. Class 3	Appl. Class 4
A	-	1 locate by key, 2 SEQ reads	1 locate by key, 5 SEQ reads	1 locate by key, 5 SEQ reads
B	-	3 Fkey reads, 1 update	11 Fkey reads, 6 updates	19 Fkey reads, 10 updates
C	1 add	1 add	1 add	1 add
D	1 Fkey read	-	-	3 Fkey reads
E	-	-	-	3 Fkey reads

Note:

- SEQ - sequential
- Fkey- full key

Table 2. RAMP-C Application and Data Relationships

5.4 System Setup

This section covers the setup of the VSE/ESA 2.1 environment including:

- VSE/ESA setup
- VSE/VTAM setup
- CICS/VSE setup
- Application and data setup

Examples of different setup jobs are provided in Appendix B, “Setup Examples” on page 49.

5.4.1 VSE/ESA Setup

The VSE/ESA 2.1 system was installed with supervisor PTF level DY43757. The following optional products were installed:

- LE/VSE Base Release 1.1.0
- LE/VSE COBOL Base Release 1.1.0
- LE/VSE COBOL CICS Release 1.1.0
- COBOL/VSE Base Release 1.0
- VSE WORKDESK - English 6.1.0
- VSE/ESA DWF 6.1.0

The VSE/ESA 2.1 system was set up to provide sufficient virtual storage for all the test runs:

- VSIZE of 400MB
- PASIZE of 80MB
- 2MB of 24-bit System GETVIS
- 8MB of 31-bit System GETVIS
- 4MB of 31-bit SVA PSIZE
- Two 75MB partitions for CICS
- SYSDEF statement with 20MB data space in the ALLOC PROC

Figure 11 provides the layout of the virtual storage map:

```
SPACE AREA      V-SIZE   GETVIS   V-ADDR   UNUSED NAME
  S   SUP         624K                 0          $$A$SUPX
  S   SVA-24     1548K    2884K    9C000       0K
  0   BG V       1280K     256K   500000   59904K
  1   F1 V        768K     832K   500000       0K POWSTART
  2   F2 V       3072K    7168K   500000       0K CICSICCF
  3   F3 V        600K   19880K   500000       0K VTAMSTRT
  4   F4 V       8192K   68608K   500000       0K
  5   F5 V       8192K   68608K   500000       0K
  6   F6 V        256K     256K   500000       0K
  7   F7 V        256K     256K   500000       0K
  8   F8 V       4096K   72704K   500000       0K
  9   F9 V       3072K   73728K   500000       0K
  A   FA V        256K     256K   500000       0K
  B   FB V        256K     256K   500000       0K
  S   SVA-31     4932K    9404K  4000000
      DYN-PA        0K
      DSPACE     3552K
      SYSTEM     1600K
      AVAIL     41952K
      TOTAL    409600K   <----¢
```

Figure 11. VSE Virtual Storage Map

5.4.2 VSE/VTAM Setup

To support the network configuration, the following VTAM customization was performed in the VSE/ESA 2.1 environment:

- VTAM startup job with 6MB of DSPACE in the EXEC statement

  Data space is required by VTAM 4.2 to communicate with VTAM applications such as CICS and PNET. The maximum data space allowed for the VTAM application is defined by the DSPACE parameter in the corresponding EXEC statement. A maximum of 6MB data space was defined in the VTAM EXEC statement and was sufficient for all the test runs.

- PATH table for the paths between VSE/VTAM, the NCP of the IBM 3745, and the VM/VTAM
- CTC connections between the two VTAMs in a VTAM CA major node
- Cross Domain Resource Manager for the VM/VTAM in a VTAM CRDM major node
- Remote LUs owned by the VM/VTAM in a VTAM CDRSC major node

5.4.3 CICS/VSE Setup

The CICS/VSE setup was changed as follows:

- CICS startup jobs with 4MB of DSPACE in the EXEC statement

  This maximum size of data space was defined to each CICS for the communication between CICS and VTAM, and was sufficient for the test runs.

- PCT and PPT were customized to include the RAMP-C transactions and programs
- FCT entries were added for the RAMP-C VSAM files
- No terminal definition for the remote LUs as the CICS auto-install function was employed
- Sufficient 31-bit partition GETVIS was allocated for the RAMP-C application programs

5.4.4 Application and Data Setup

The following preparation was performed for the RAMP-C applications and data:

- 31-bit SVA eligible phases for LE/VSE 1.1 were placed in the 31-bit SVA
- the RAMP-C application programs to support the VSAM files were compiled and cataloged to different user sublibraries
- the VSAM files were defined and loaded to different VSAM user catalogs

Chapter 6. Test Cases

The objective of this chapter is to show the performance effects of different system parameters on the utilization of a multiprocessor using the VSE/ESA 2.1 Turbo Dispatcher.

This chapter consists of the following sections:

- Test Plan - describes the design and logic of the test runs
- Measurement Readings - describes the data collected for the test results
- Test Runs - discusses the test runs and findings
- Test Runs Summary - summarizes the test results

6.1 Test Plan

The key element affecting the utilization of a multiprocessor for a certain workload, is the ratio of CPU time consumed by the non-parallel work units to the total CPU time (**Non-Parallel Share**).

A lower non-parallel share means:

- Less non-parallel work units to be processed by the system
- Less wait time for the partitions to become non-parallel
- More parallel work units to be processed
- More Central Processors (**CP**) to be exploited

There are two key areas in which the VSE/ESA 2.1 user can affect the non-parallel share:

- Increase the number of concurrently dispatchable parallel work units by increasing the number of concurrent active partitions
- Reduce the number of non-parallel work units by reducing the number of physical I/Os

The following test runs were designed:

- Single CP with the Standard Dispatcher - to determine the base. This test run can be regarded as the direct migration step from a VSE/ESA V.1 to a VSE/ESA 2.1 environment, without using the multiprocessor support.
- Single CP with the Turbo Dispatcher - to determine the system resource requirements of the Turbo Dispatcher. This test run can be regarded as the migration step from a VSE/ESA V.1 to a VSE/ESA 2.1 environment, with the Turbo Dispatcher.
- Multiprocessor with the Turbo Dispatcher - to determine the performance effects of adding CPs to a CPU constrained environment. This test run shows a VSE/ESA 2.1 environment with the multiprocessor utilization.
- Multiple CICS partitions - to determine the performance effects of more concurrently dispatchable partitions with the Turbo Dispatcher. This test run can be regarded as the splitting of two CICS partitions into three functionally independent CICS partitions.

- S/390 Hardware-Assisted Data Compression - to determine the performance effects when reducing physical DASD I/Os by using the S/390 Hardware-Assisted Data Compression with VSAM files.

The test runs used the RAMP-C workload with VSAM, as the majority of VSE/ESA installations have VSAM as their preferred file access method.

6.2 Measurement Readings

After the workload was stabilized the duration of each test run was 15 minutes. For each test run, the performance data from various sources was gathered and documented in a summary table for comparison:

- Transaction throughput measured by tasks per minute (**Tasks/Min**) from the TPNS response statistics report; one reported TPNS task per minute can be seen as one CICS transaction per minute.
- Mean response time in seconds (**Resp. Time**) from the TPNS response statistics report.
- CPU utilization for the entire VSE/ESA system and the active partitions (**CPU%**) from the Interactive Interface (II) Display System Activity panel.
- Non-parallel share (**NPS**) from the QUERY TD command.
- I/Os per second (**IOs/Sec**) from the II Display System Activity panel. The I/O rate from the II panel includes all the I/Os in the system during the sampling period. It contains disk, terminal and other I/Os.

During all the test runs, the following system resources were also monitored:

- Partition and system GETVIS using the appropriate GETVIS commands
- CPU utilization and paging rate, using the II Display System Activity panel

No paging activity was observed for any of the test runs.

The tracking of these major system resources was to ensure that the maximum load was reached, without any performance bottleneck except that of the CPU.

Disclaimer

Any performance data contained in this document was determined in a **controlled and unconstrained** environment, and therefore, the results that may be obtained in other operating environments may vary significantly.

6.3 Basic Test for the Standard Dispatcher

This section describes the test runs to establish the appropriate workload running under the Standard Dispatcher, to be compared with the Turbo Dispatcher test runs in the next sections. The test runs also represent a direct migration step from a VSE/ESA V.1 to a VSE/ESA 2.1 environment using the Standard Dispatcher in both releases.

The objective was to determine the maximum load which puts the single CP under pressure, without introducing other performance bottlenecks. The following workloads were tested with the Standard Dispatcher:

- RAMP-C with VSAM - One CICS partition with 350 LUs
- RAMP-C with VSAM - Two CICS partitions with 350 LUs per CICS partition

The additional CICS partition in the second test run had the same priority and characteristics as the first CICS partition, with VSAM files defined in a different user catalog and located on different DASD volumes. Adding the workload of one CICS with the same number of active LUs implied that twice the workload was offered to the system.

6.3.1 Test Results

The test results are summarized in the tables below:

Test Run	Tasks/ Min	Resp. Time	CPU %	CPU % F5/F4	NPS	IOs/Sec
1-CICS w 350 LUs	1437	0.74	61 %	61 %	-	168
2-CICS w 350 LUs	2342	5.69	99 %	50% 49%	-	260
Note: NPS - non-parallel share						

Table 3. Performance Results - Standard Dispatcher

The GETVIS utilization is summarized below:

Test Run	GETVIS	Allocated 24-bit	Max-used 24-bit	Allocated 31-bit	Max-used 31-bit
1-CICS w 350 LUs	CICS F5	2048K	448K	64512K	60964K
	System	1716K	1328K	9400K	3352K
2-CICS w 350 LUs	CICS F4	3072K	472K	64476K	60960K
	CICS F5	3072K	484K	64476K	62732K
	System	1716K	1364K	9400K	3544K

Table 4. GETVIS Utilization - Standard Dispatcher

6.3.2 Observations

As the workload increased, the transaction throughput, response time, CPU utilization for the system and partitions, and I/O rate increased.

The workload of ″**Two CICS with 350 LUs per CICS**″ reflected a totally CPU constrained environment with high transaction response times. This amount of work was therefore selected to be compared with the Turbo Dispatcher test runs in the next section.

61MB of 31-bit GETVIS was used in each CICS partition. It was mainly used by:

- Large and multiple VSAM LSR buffer pools
- 31-bit eligible LE/VSE code
- 31-bit eligible application programs

The 24-bit and 31-bit system and partition GETVIS areas were sufficient for both loads.

6.4 Standard Dispatcher to Turbo Dispatcher Transition

The objective of this test run was to determine the system resource requirements of the Turbo Dispatcher.

The test run in this section represents an important transition step from the single CP to a multiprocessor environment. Similar test runs with a customer′s workload should be conducted to evaluate:

- The performance effects on the existing workload when using the Turbo Dispatcher
- The non-parallel share of the workload to determine if the workload is applicable in a multiprocessor environment

With the Turbo Dispatcher activated for the same environment, a slight increase in the CPU utilization and system GETVIS is anticipated. It is mainly due to the increase in functions for the Turbo Dispatcher. Overall system performance should be similar to the Standard Dispatcher results for the same workload.

Using the RAMP-C with VSAM workload of ″Two CICS with 350 LUs per CICS″ from the previous section, the same test run was conducted with the Turbo Dispatcher.

The performance data and the GETVIS utilization were recorded for comparison with the Standard Dispatcher test results.

6.4.1 Test Results

The test results are summarized in the tables below:

Test Run **2 CICS/350 LUs each**	**Tasks/ Min**	**Resp. Time**	**CPU %**	**CPU % F5/F4**	**NPS**	**IOs/Sec**
SD with 1 CP	2342	5.69	99 %	50% 49%	-	260
TD with 1 CP	2270	4.95	99 %	50% 49%	0.499	262

Note:

- SD - Standard Dispatcher
- TD - Turbo Dispatcher
- NPS - non-parallel share

Table 5. Performance Results - Turbo Dispatcher

Test Run	**GETVIS**	**Allocated 24-bit**	**Max-used 24-bit**	**Allocated 31-bit**	**Max-used 31-bit**
SD with 1 CP	CICS F4	3072K	472K	64476K	60960K
	CICS F5	3072K	484K	64476K	62732K
	System	1716K	1364K	9400K	3544K
TD with 1 CP	CICS F4	3072K	468K	64476K	60908K
	CICS F5	3072K	444K	64476K	60544K
	System	1688K	1388K	9404K	3640K

Table 6. GETVIS Utilization - Turbo Dispatcher

6.4.2 Observations

The system requirements of the Turbo Dispatcher are minimal, based on the test results:

- CPU utilization naturally remained at 99%.
- The transaction throughput decreased slightly from 2342 to 2270 tasks per minute.
- The transaction response time decreased from 5.69 to 4.95 seconds, due to the reduction of transaction throughput.
- I/Os per second were stable.
- A slight increase (96K) in the used 31-bit system GETVIS was observed. This was due to the requirements for the Turbo Dispatcher work areas and control blocks.

Using the QUERY TD command, a non-parallel share of 0.499 was observed for this workload. Based on Table 1 on page 8, this particular workload could take advantage of the VSE/ESA 2.1 multiprocessor support. The non-parallel share of this workload suggested that up to 1.8 CPs may be fully exploited, if sufficient partitions could be established.

Multiprocessor Planning

Determining the non-parallel share of the user specific workload is an important step in the processor capacity planning. Testing the customer's workload under the Turbo Dispatcher with a single CP will provide the user with an indication of whether the workload may have advantages running in a multiprocessor environment. It also assists the user to estimate the maximum number of CPs that the workload may exploit at best.

The non-parallel share changes as the system environment and workload changes. However, an initial assessment will give the user a frame of reference for processor capacity planning purposes.

6.4.3 Conclusions

The following conclusions can be drawn, based on the test results:

- The Turbo Dispatcher system resource requirements in terms of CPU utilization and system GETVIS are minimal.
- Running the workload using the Turbo Dispatcher with a single CP provided the non-parallel share required for the processor capacity planning. The non-parallel share in the test run suggests that up to 1.8 CPs may be effectively exploited by this workload.

6.5 Turbo Dispatcher on a Multiprocessor

The objective of this test run was to determine the performance effects when additional CPs are introduced into a CPU constrained environment.

The test results of the previous section reflect a CPU constrained environment. The non-parallel share of the workload indicate that an additional CP could be exploited. When an additional CP is made available to the same workload, the performance bottleneck indicated by the high CPU utilization should be removed.

As a result, improvements in transaction throughput and response time were anticipated.

The RAMP-C with VSAM workload and the environment of the previous section were used for this test run. The only system variable introduced was the number of CPs. Two runs were conducted; one with two CPs and one with three CPs. The purpose of the second test run was to test the system behavior when more CPs were made available than the non-parallel share suggested.

The performance data and GETVIS utilization were recorded for comparison with the single CP test run.

6.5.1 Test Results

The test results are summarized in the tables below:

Test Run **2 CICS/350 LUs each**	**Tasks/ Min**	**Resp. Time**	**CPU %**	**CPU % F5/F4**	**NPS**	**IOs/Sec**
TD with 1-CP	2270	4.95	99 %	50% 49%	0.499	262
TD with 2-CP	2675	2.25	125 %	63% 62%	0.533	297
TD with 3-CP	2630	2.37	123 %	62% 61%	0.533	293
Note: NPS - non-parallel share						

Table 7. Performance Results - Turbo Dispatcher and Multiprocessor

Test Run	**GETVIS**	**Allocated 24-bit**	**Max-used 24-bit**	**Allocated 31-bit**	**Max-used 31-bit**
TD with 1-CP	CICS F4	3072K	468K	64476K	60908K
	CICS F5	3072K	444K	64476K	60544K
	System	1688K	1388K	9404K	3640K
TD with 2-CP	CICS F4	3072K	460K	64476K	63312K
	CICS F5	3072K	456K	64476K	58222K
	System	1688K	1392K	9404K	3724K
TD with 3-CP	CICS F4	3072K	448K	64476K	60888K
	CICS F5	3072K	436K	64476K	60188K
	System	1688K	1472K	9404K	3588K

Table 8. GETVIS Utilization - Turbo Dispatcher and Multiprocessor

6.5.2 Observations - One and Two CPs

Compared with the results of the single CP test run, the two-CP test run produced impressive performance improvements. When the second CP was introduced to the workload:

- The transaction throughput increased from 2270 to 2675 tasks per minute.
- The transaction response time decreased from 4.95 to 2.25 seconds. On average, the transactions ran 2.2 times faster.

- CPU utilization increased to 125%, and spread evenly between the two CICS partitions.
- I/Os per second increased from 262 to 297, as a result of the increased transaction throughput.
- A slight increase (84K) in the 31-bit system GETVIS utilization for the Turbo Dispatcher work areas and control blocks was observed.
- The non-parallel share increased from 0.499 to 0.533, as a result of the additional non-parallel work units to manage the additional CP and the increased I/O counts. The new non-parallel share of 0.533 suggested that up to 1.7 CPs could be exploited by this workload.

6.5.3 Observations - Two and Three CPs

Comparing the test results of the two-CP and three-CP test runs, illustrated the effects of adding an additional CP to a CPU unconstrained system.

The basic rule of performance tuning is to identify and remove the bottlenecks that cause performance problems. After the second CP was added to the system, the CPU utilization of the system was averaged at 125%. The performance bottleneck was no longer the available total CPU power.

Additionally, the non-parallel share of 0.533 suggested that this particular workload may only exploit up to 1.7 CPs. Therefore, adding another CP to this workload should yield no performance benefit.

Compared with the test results of the two-CP test run, the third CP provided the following performance effects:

- The transaction throughput decreased from 2675 to 2630 tasks per minute.
- The transaction response time increased from 2.25 to 2.37 seconds.
- CPU utilization was reduced to 123% and spread evenly between the two CICS partitions. An indication that the additional CP was not used effectively.
- I/Os per second decreased slightly as a result of the reduced transaction volume.
- The non-parallel share remained at 0.533, indicating that the third CP could not be utilized effectively.

6.5.4 Conclusions

When adding the appropriate number of CPs to a CPU constrained environment, the additional investment pays off easily by the performance benefits achieved in the system. The performance benefits may be realized in terms of:

- Improvement in transaction throughput
- Reduction in transaction response time

The test results of the two-CP environment demonstrate how the VSE/ESA 2.1 user can achieve significant performance gains by adding a CP to a CPU constrained environment.

The current Turbo Dispatcher implementation does not allow a partition to exploit more than one CP at any one time. Adding an unnecessary CP to an unconstrained environment will not provide any incremental performance benefit. In certain situations, it will even bring inverse performance effects to the VSE/ESA 2.1 environment, mainly due to the resource requirements to manage the additional CP.

Careful processor configuration planning is therefore necessary for the VSE/ESA 2.1 multiprocessor environment.

6.6 Multiple CICS Partitions

The objective of the test runs was to illustrate the performance effects of splitting CICS partitions into multiple functionally independent CICS partitions in a multiprocessor environment.

The non-parallel share may be reduced by increasing the number of parallel work units in the system. More parallel work units can be introduced to the system by adding partitions with a workload consisting of a lower non-parallel share. If sufficient CPU cycles and CPs are available, the non-parallel share should be reduced and a better multiprocessor utilization is anticipated. It should result in a better overall system performance, normally reflected by:

- Increased transaction throughput
- Better transaction response times

Splitting CICS partitions into multiple functionally independent CICS partitions produces more concurrent active partitions in the system. If real storage is available for the additional virtual storage requirements of the multiple CICS partitions, further performance benefits may be realized as a result of:

- More Dynamic Storage Area (DSA)

  Fewer terminals per CICS partition imply a reduction of DSA requirements, as fewer Terminal Control Table Terminal Entries (TCTTE) and related control blocks are needed.

  Fewer transactions per CICS imply a reduction of Task Control and Program Control storage requirements per transaction.

- More available partition GETVIS

  Fewer transactions per CICS imply fewer file I/Os could be made. It implies that more buffers in the LSR pools are available for the active transactions.

The combined performance effects should be taken into consideration when evaluating the test results.

The workload and test results of ″Two CICS with 350 LUs per CICS under two-CPs″ of the previous section were chosen as the base. The test run in this section used the RAMP-C with VSAM workload under the following environment:

- Three CICS partitions with the same priority
- 250 LUs per CICS
- Two CPs

This workload and environment were designed to simulate the splitting of the two CICS partitions into three functionally independent CICS partitions.

The additional CICS partition in this test run had the same priority and characteristics as the other CICS partitions, but with VSAM files defined in a different user catalog and located on different DASD volumes.

A total of 750 LUs in this test run represented a 7% increase in workload when compared with 700 LUs in the base workload. The slight workload increase was designed for the small terminal and transaction duplications, normally accrued when splitting a CICS partition into multiple functionally independent CICS partitions. The slight workload increase was taken into consideration when evaluating the test results.

The performance data and GETVIS utilization were recorded for comparison with the former ″Two CICS with two-CP″ test runs.

6.6.1 Test Results

The test results are summarized in the tables below:

Test Run	Tasks/ Min	Resp. Time	CPU %	CPU % F5/F4/ F9	NPS	IOs/Sec
2-CICS with 2-CP	2675	2.25	125 %	63% 62%	0.533	297
3-CICS with 2-CP	3286	0.51	126 %	43% 42% 41%	0.462	387

Note:

All CICS partitions carry the same characteristics with files defined to a different user catalog and located in different volumes

NPS - non-parallel share

Table 9. Performance Results - Multiple CICS Partitions

Test Run	GETVIS	Allocated 24-bit	Max-used 24-bit	Allocated 31-bit	Max-used 31-bit
2-CICS with 2-CP 350 LUs each	CICS F4	3072K	460K	64476K	63312K
	CICS F5	3072K	456K	64476K	58222K
	System	1688K	1392K	9404K	3724K
3-CICS with 2-CP 250 LUs each	CICS F4	3072K	432K	64476K	52328K
	CICS F5	3072K	432K	64476K	52328K
	CICS F9	3072K	428K	44768K	53644K
	System	1688K	1440K	9404K	3796K

Table 10. GETVIS Utilization - Multiple CICS Partitions

6.6.2 Observations

When evaluating the test results, the following should be taken into consideration:

- Possible performance improvement due to the reduction of the non-parallel share
- Possible performance improvement due to the availability of DSA and partition GETVIS

- Performance effects due to the increased workload from 700 to 750 LUs

Compared with the result of the two-CICS test run, the three-CICS test run produced the following performance improvements:

- The transaction throughput increased from 2675 to 3286 tasks per minute.

  To compare the two transaction throughputs, the transaction throughput per LU was obtained by dividing the transactions throughput with the number of LUs.

  - 2675 / 700 = 3.8 tasks/min/LU for the 700 LU workload
  - 3286 / 750 = 4.3 tasks/min/LU for the 750 LU workload

  The weighted result demonstrated that the three-CICS test run provided a much higher transaction throughput.

- The transaction response time decreased significantly from 2.25 to 0.51 seconds. On average, the transactions ran 4.4 times faster.
- CPU utilization was almost unchanged, and spread evenly among the three CICS partitions. The significant improvements in the transaction performance indicated that the CPU cycles were utilized more productively.
- I/O per second increased from 297 to 387, as a result of the increased transaction throughput.
- The non-parallel share decreased from 0.533 to 0.462 as result of more parallel work units in the system.

6.6.3 Conclusions

The test run can be regarded as the splitting of two CICS partitions into three functionally independent CICS partitions.

The test results demonstrate that this approach provides a good performance tuning alternative in the VSE/ESA 2.1 multiprocessor environment, as:

- Introducing more parallel work units using multiple CICS partitions allows the VSE/ESA 2.1 Turbo Dispatcher to better utilize the available CPs.
- Splitting large CICS partitions into multiple CICS partitions without changing the partition virtual storage layout, provides more CICS resources such as 24-bit DSA and GETVIS, to each individual transaction.

As a result, the test run produced an impressive overall system performance benefit with improved transaction throughput and response time.

6.7 S/390 Hardware-Assisted Data Compression

S/390 Hardware-Assisted Data Compression allows the user to utilize this hardware function to perform data compression, instead of the VSAM Software Data Compression, with the following ESA/390 processors:

- IBM ES/9000 Processor 711-based models
- IBM ES/9000 Processor 511-based models
- IBM ES/9000 Processor 211-based models
- IBM S/390 9672 Parallel Enterprise Servers

Both S/390 Hardware-Assisted Data Compression and VSAM Software S/390 Hardware-Assisted Data Compression support VSAM ESDS, KSDS and VRDS files. The compression and expansion are done on the record level.

The objective of the test runs was to determine the performance effects when S/390 Hardware-Assisted Data Compression was used for the VSAM files.

The immediate benefit of the data compression is the reduction of disk space required. A more important benefit from the system performance aspect is the potential reduction of physical I/Os.

With a good compression factor, more logical records can be stored and retrieved in one Control Interval (CI). If the number of buffers in the LSR pools remains unchanged, more data can be stored in memory and a reduction of physical I/Os could be expected.

The reduction of physical I/Os should reduce the non-parallel share of the workload. This should result in an overall improvement of the system performance.

For the test runs, all the VSAM files, with the exception of the application log C files, were redefined with the COMPRESSED parameter in the IDCAMS DEFINE CLUSTER commands. The files were reloaded and the data dictionaries for the data compression were created.

All the files were used to run with the RAMP-C with VSAM workload and ″**Two CICS with 350 LUs per CICS**″ under two CPs. Based on the test results, another run was conducted with three CPs. The results were compared with the test run without the use of the S/390 Hardware-Assisted Data Compression.

Data Dictionary

Data compression relies on the data dictionary to perform the compression and expansion functions. This data dictionary is built when the file is initially loaded.

C files are application log files. They were first initialized and then opened for update for each test run. Therefore, the data dictionary could not be built for C files, and they remained as normal VSAM data sets.

Compression Factor

The compression factor was obtained by comparing the difference in the space utilization of the same file, before and after the file was compressed. The VSAM LISTCAT of all the files were produced, and the space utilizations of the files were calculated using data from:

- Allocation type and unit
- Number of primary and secondary extents
- High-Allocated-RBA and High-Used-RBA

The compression factor of each file was computed by:

compression factor = decompressed space / compressed space

The average compression factor of the test files was 1.37. That implied 37% more compressed logical records could be stored and retrieved in the same CI.

6.7.1 Test Results

The test results are summarized in the tables below:

Test Run 2 CICS/350 LUs each	Tasks/ Min	Resp. Time	CPU %	CPU % F4/F5	NPS	IOs/Sec
2-CICS with 2-CP	2675	2.25	125 %	63% 62%	0.533	297
2-CICS, 2-CP with Compression	3032	0.67	142 %	73% 69%	0.415	287
2-CICS, 3-CP with Compression	3041	0.64	151 %	76% 75%	0.417	291
Note: NPS - non-parallel share						

Table 11. Performance Results - S/390 Hardware-Assisted Data Compression

Test Run	GETVIS	Allocated 24-bit	Max-used 24-bit	Allocated 31-bit	Max-used 31-bit
2-CICS with 2-CP	CICS F4	3072K	460K	64476K	63312K
	CICS F5	3072K	456K	64476K	58222K
	System	1688K	1392K	9404K	3724K
2-CICS, 2-CP with Compression	CICS F4	1024K	440K	65536K	61116K
	CICS F5	1024K	444K	65536K	59472K
	System	2712K	1388K	9404K	3684K
2-CICS, 3-CP with Compression	CICS F4	1024K	452K	65536K	58796K
	CICS F5	1024K	436K	65536K	59500K
	System	2712K	1392K	9404K	3736K

Table 12. GETVIS Utilization - S/390 Hardware-Assisted Data Compression

6.7.2 Observations - Two CPs

Comparing the test results of the two runs, the S/390 Hardware-Assisted Data Compression test run provided the following performance data:

- The transaction throughput increased from 2675 to 3032 tasks per minute.
- The transaction response time decreased significantly from 2.25 to 0.67 seconds. On average, the transactions ran 3.4 times faster.
- CPU utilization increased from 125 to 142%, spread evenly among the two CICS partitions. The overall improvements in transaction performance suggested that the CPU overhead of S/390 Hardware-Assisted Data Compression was insignificant.
- I/Os per second decreased from 297 to 287.
- The non-parallel share decreased from 0.533 to 0.415. A non-parallel share of 0.415 implied that up to 2.1 CPs could be exploited by this workload with the compressed VSAM files.

6.7.3 Observations - Two and Three CPs

Compared with the test result of the two-CP test run, the third CP provided the following performance effects:

- The transaction throughput was almost unchanged.
- The transaction response time decreased slightly from 0.67 to 0.64 seconds.
- CPU utilization increased from 142 to 151%, mainly due to the CPU cycles required by the Turbo Dispatcher to manage the additional CP.
- I/Os per second were almost unchanged.
- The non-parallel share increased slightly from 0.415 to 0.417.

6.7.4 Conclusions

The usage of S/390 Hardware-Assisted Data Compression for VSAM files produces very positive performance results when CPU cycles are available for the function. It is an attractive tuning alternative for the ESA/390 user, as the function:

- Produces very positive performance results
- Is easy to implement
- Requires no changes to applications
- Reduces DASD and buffer space requirements

The third CP did not provide any performance benefit for this workload. It conforms to the non-parallel share suggestion that this workload could not effectively exploit three CPs. This shows again that careful processor capacity planning is needed to warrant a reasonable return on investment for the additional CP in a VSE/ESA 2.1 multiprocessor environment.

6.8 Test Runs Summary

The objective of this chapter was to demonstrate the performance effects of different system parameters on the utilization of a multiprocessor using the VSE/ESA 2.1 Turbo Dispatcher. To achieve this objective, test runs were designed and performed to illustrate:

- The exploitation possibilities of a multiprocessor in a VSE/ESA 2.1 environment
- The improvement of the multiprocessor utilization using various system facilities

The test results in chapter 6.5, "Turbo Dispatcher on a Multiprocessor" on page 33, demonstrate how a VSE/ESA 2.1 user could improve the overall system performance by taking advantage of the multiprocessor support of VSE/ESA 2.1. The test results also indicate that careful planning of the processor capacity in a VSE/ESA 2.1 multiprocessor environment is needed, as an unnecessary number of CPs may be counter-productive.

The test results in chapter 6.6, "Multiple CICS Partitions" on page 36 illustrate the potential performance benefits of splitting a single CICS partition into multiple functionally independent CICS partitions in a VSE/ESA 2.1 multiprocessor environment.

The test results in chapter 6.7, "S/390 Hardware-Assisted Data Compression" on page 38 illustrate the insignificant cost involved to utilize the S/390 Hardware-Assisted Data Compression and the resulting performance benefits.

In summary, all the test runs performed in this chapter deliver promising performance results in terms of improved transaction throughput and response time. The results also demonstrate how the VSE/ESA 2.1 user may better utilize a multiprocessor in a VSE/ESA 2.1 environment.

Chapter 7. Guidelines for Using the VSE/ESA Turbo Dispatcher

The objective of this chapter is to provide guidelines for using the VSE/ESA 2.1 Turbo Dispatcher in a multiprocessor environment, based on the results from the various performance test runs described in Chapter 6, “Test Cases” on page 29.

This chapter consists of the following sections:

- Multiprocessor Considerations - describes the considerations when sizing a multiprocessor system for the VSE/ESA 2.1 user
- Migration Considerations - suggests migration steps as related to the VSE/ESA 2.1 Turbo Dispatcher in a multiprocessor system
- Partition Considerations - shows potential options to better utilize the Turbo Dispatcher by introducing additional work to the system through a higher number of concurrent active partitions
- Data In Memory (DIM) Considerations - considers the usage of DIM facilities available to the VSE/ESA 2.1 user to decrease the non-parallel share of a workload

7.1 Multiprocessor Considerations

When sizing a multiprocessor for the VSE/ESA 2.1 user, the following factors must be considered:

- The maximum utilization of one CP

  The Turbo Dispatcher dispatches units of work on the partition level. Therefore, a single partition can be active only on one CP at a time. Thus one CP is the maximum capacity that one partition may exploit.

- CPU utilization of the partitions

  When sizing the capacity of a multiprocessor system, the requirement of the biggest CPU consuming partition must be considered.

  Consider a uniprocessor system with a single CICS partition being the biggest CPU cycle consumer. CPU utilization of this partition averages 80%. When upgrading this system to a multiprocessor system, at least a CP with an equivalent capacity must be available.

  A multiprocessor system with less capacity per CP may be configured only if the biggest CPU consuming partition can be divided into multiple partitions similar to the example in chapter 6.6, “Multiple CICS Partitions” on page 36. Note that the biggest of the multiple CICSs still must obey this rule.

- Active partitions and the number of CPs

  The test results in Chapter 6, “Test Cases” on page 29 demonstrate that:

  - Adding concurrent active partitions could bring increased overall throughput to an unconstrained environment.
  - Adding unnecessary CPs to a system will not bring any performance benefit, but rather may degrade the overall CPU time per transaction.

  Based on these observations, a multiprocessor system must have enough concurrently dispatchable partitions to exploit the aggregate CPU power.

Regarding the CPU requirements, the following must also be considered:

- CPU time for the Turbo Dispatcher

  5-10% of the CPU should be added to the system for the additional Turbo Dispatcher processing functions.

- Customer requirements and growth

  As well as all the above considerations, the customer data processing requirements and growth rate remain as the key decision criteria to determine the CPU capacity, in terms of:

  – The processing capacity of one CP

  – The total number of CPs required

7.2 Migration Considerations

The test runs with the Standard Dispatcher and Turbo Dispatcher in Chapter 6, "Test Cases" on page 29 provide hints for the VSE/ESA 2.1 user migrating from a uniprocessor to a multiprocessor environment. The following migration steps should be taken:

- Determine the non-parallel share of the current workload

  This step can be done with the existing processor as the first migration step.

  This is similar to the test run in chapter 6.4, "Standard Dispatcher to Turbo Dispatcher Transition" on page 32, to determine the non-parallel share of the workload. Online and batch environments may carry different processing characteristics, therefore different values of non-parallel share are expected. Specific vendor or IBM programs will influence the non-parallel share of a workload as well.

- Activate one CP at a time

  Based on the non-parallel share of the current workload, the maximum number of CPs that the workload may exploit at best can be estimated.

  The testing should be repeated in the installed multiprocessor system, with one additional CP activated at a time, up to the estimated number of CPs. A minimal increase of the non-parallel share will be seen with the increasing number of CPs used.

  The system performance indicators such as online transaction throughput, response time, elapsed time for batch jobs and non-parallel share must be monitored, to ensure that the system achieves the customer′s performance expectation.

- Reduce the non-parallel share

  After the non-parallel share of the workload under the multiprocessor environment is obtained, further tuning should be performed to reduce the non-parallel share. The target is to allow the workload to better exploit the installed multiprocessor system and to get better online response times or batch elapsed times.

7.3 Partition Considerations

The test results in Chapter 6, “Test Cases” on page 29 indicate that an option to better utilize the Turbo Dispatcher is to introduce additional work through additional active partitions. As a result, an overall system performance improvement was realized in an unconstrained multiprocessor system. The applicability of this approach should be evaluated for the customer′s workload, for example by:

- Using more concurrent batch jobs

  Determine if it is possible to run more batch jobs concurrently with the online environment. This may help to reduce the non-parallel share and to improve the batch window. The performance indicators must be monitored to ensure that the online environment maintains an acceptable performance level. It must be observed as to whether the additional concurrent batch work impacts the online production work in terms of locks and I/O contention.

- Splitting a large CICS partition

  Splitting a large CICS partition into multiple CICS partitions involves a technical feasibility study, planning, implementation and testing. With the appropriate system resources available, the effort could be justified by the expected performance benefits.

  There are different approaches to divide a large CICS partition into multiple CICS partitions, for example:

  – Split by functions or applications

  – Divide by Multiregion Operation (MRO) with a combination of Terminal Owning Regions (TORs), Application Owning Regions (AORs) and File Owning Regions (FORs)

- Using multiple SQL/DS 3.5 application servers

  The applicability of the SQL/DS 3.5 Dynamic Database Switching facility should be evaluated for the customer′s workload.

  The optimal number of SQL/DS application servers depends mainly on the customer′s database usage and the availability of system resources. If applicable using two SQL/DS application servers could be a good starting point. The system performance indicators must be closely monitored for each additional SQL/DS application server partition introduced.

  The VSE/ESA 2.1 user may take advantage of this facility to distribute the load of a single SQL/DS application server across multiple SQL/DS application servers to better exploit the Turbo Dispatcher.

  Therefore, distributing the load to multiple SQL/DS application servers in a VSE/ESA 2.1 environment provides another tuning potential on multiprocessor systems.

7.4 Data In Memory Considerations

The main option to decrease the non-parallel share of a workload and thus to increase the system performance is to place more data into storage by using Data In Memory (DIM) facilities or the S/390 Hardware-Assisted Data Compression.

More data can be placed in storage by using software DIM facilities, for example:

- VSAM buffers in 31-bit GETVIS

  VSAM places its data buffers in 31-bit partition GETVIS, if available. Sufficient allocation of GETVIS can be provided by increasing the size of the partition (ALLOC).

  Since VSE/ESA 1.3, there are 15 VSAM LSR pools available which simplifies the control and tuning of VSAM buffers. CICS statistics provide the utilization of the LSR pools with buffer hit data.

  LSR is the preferred method for files with mostly random accesses. NSR requires more virtual storage, and is more suitable for files with mostly sequential access.

- CICS Data Tables

  This facility was first introduced in CICS/VSE 2.2 with VSE/ESA 1.3. It allows the user to place the most active KSDS files, wholly or partly, in 31-bit partition GETVIS.

  There are two types of CICS data tables: CICS Maintained Data Tables (CMT) and User Maintained Data Tables (UMT). Using CMTs requires no change in the application, while UMTs provide new application design opportunities to the user.

- Virtual Disk

  Virtual Disk was first introduced in VSE/ESA 1.3, allowing the user to define virtual FBA disks in data space. The usage of Virtual Disks is mostly effective for work files during batch processing.

  Caution: The processing of Virtual Disk is mainly non-parallel. The usage of Virtual Disk will therefore increase the non-parallel share, and thus may affect multiprocessor utilization in a VSE/ESA 2.1 environment.

Virtual to Central Storage Ratio

Software DIM implementation requires virtual storage. Sufficient central storage must be available to support the additional virtual storage required by the DIM facilities used. It is counter-productive to trade file I/Os for paging I/Os.

As demonstrated in the test results in chapter 6.7, "S/390 Hardware-Assisted Data Compression" on page 38, system performance improvement can be achieved by placing more data in storage by using S/390 Hardware-Assisted Data Compression.

S/390 Hardware-Assisted Data Compression has the following attractions:

- It is easy to implement. It involves file unload, DELETE/DEFINE and reload. It could result in a quick performance return with minimal effort.
- The CPU cycle requirement for S/390 Hardware-Assisted Data Compression is relatively small compared to the equivalent software data compression.

The applicability of this hardware function should be evaluated, especially for heavy I/O workloads with good data reference locality. To estimate the applicability of S/390 Hardware-Assisted Data Compression, the tool IKQCPRED from the IBM tools disk IBMVSE can be used.

Appendix A. IPL Load Parameter

During IPL, the Turbo Dispatcher is loaded into shared storage above the supervisor ready to take over control for dispatching.

The IPL load parameter can be used to request:

1. A preferred system console device.
2. Suppression of IPL messages and command logging.
3. Prompting for IPL parameters.
4. Prompting for system startup mode.
5. Activation of the Turbo Dispatcher.

The IPL load parameter format is as follows:

```
1   2   3   4   5   6   7   8

I   S   P   P   T   X   X   X



                       Reserved

                   Turbo Dispatcher Activation

                Startup Mode Prompting

             IPL Parameter Prompting

          IPL Message Suppression

      Console Type
```

To reduce the possibility of entering wrong information, a period (.) is allowed and should be used as place-holder character. For a position that contains a period, IPL selects a default value as described below.

Parameter Description

1. The **Console Type** code specifies whether the messages are to be routed to the integrated console or to a local console. Possible values are:

   blank If the console specified in the IPL procedure is operational, the messages are routed to that local console. If the local console is not known or not available, the system waits for an interrupt from a local console. **This is the default**.

   .(PERIOD) Same as blank.

   L Route messages to a local console; this console selection is the same as for blank.

   I Route messages to the integrated console; if the integrated console is not available, the system selects a local console.

2. The **IPL Message Suppression** code can be used to request the suppression of messages and command logging during IPL. Possible values are:

   blank Display all IPL messages and print IPL commands on the system console unless the NOLOG option is specified in the supervisor parameter command. **This is the default**.

   .(PERIOD) Same as blank.

   S Suppress all informational messages during IPL. Print only error messages that require a response or an action. Do not print IPL commands on the system console.

3. The **IPL Parameter Prompting** code can be used to request a prompting for IPL parameters. Possible values are:

   blank Do not prompt for IPL parameters. **This is the default**.

   .(PERIOD) Same as blank.

   P Print message 0I03D which prompts for IPL parameters.

4. The **Startup Mode Prompting** code can be used to request a prompting for a system startup mode such as BASIC or MINI. Possible values are:

   blank Do not prompt for startup mode. **This is the default**.

   .(PERIOD) Same as blank.

   P Print messages IESI0214I and IESI0215A which prompt for the startup mode.

5. The **Turbo Dispatcher** code can be used to activate the Turbo Dispatcher instead of the Standard Dispatcher. If the Turbo Dispatcher has been successfully activated, message 0J65I is issued. Possible values are:

   blank Activate the Standard Dispatcher. **This is the default**.

   .(PERIOD) Same as blank.

   T Activate the Turbo Dispatcher.

Appendix B. Setup Examples

B.1 VSE IPL Procedure

```
089,$$A$SUPX,VSIZE=400M,VIO=512K,VPOOL=64K,LOG
ADD 00C,2540R
ADD 00D,2540P
ADD 00E,1403
ADD 080:089,3277
ADD 200:207,ECKD
ADD 840,CTCA,EML
ADD 990,3480
ADD 991,3480,00
ADD FEC,3505            POWER DUMMY READER, DO NOT DELETE
ADD FED,2520B2          POWER DUMMY PUNCH, DO NOT DELETE
ADD FEE,PRT1            POWER DUMMY PRINTER, DO NOT DELETE
ADD FEF,PRT1            POWER DUMMY PRINTER, DO NOT DELETE
ADD FFA,3505            ICCF INTERNAL READER, DO NOT DELETE
ADD FFC,3505            ICCF DUMMY READER, DO NOT DELETE
ADD FFD,2520B2          ICCF DUMMY PUNCH,  DO NOT DELETE
ADD FFE,PRT1            ICCF DUMMY PRINTER, DO NOT DELETE
ADD FFF,CONS            DUMMY CONSOLE, DO NOT DELETE
SET ZONE=WEST/00/00
DEF SYSCAT=DOSRES
DEF SYSREC=SYSWK1
SYS JA=YES
SYS BUFSIZE=1500
SYS NPARTS=44
SYS DASDFP=NO
SYS SEC=NO
SYS PASIZE=80M
SYS SPSIZE=0K
SYS BUFLD=YES
DPD VOLID=DOSRES,CYL=244,NCYL=15,TYPE=N,DSF=N
DPD VOLID=SYSWK1,CYL=501,NCYL=15,TYPE=N,DSF=N
DPD VOLID=DOSRES,CYL=465,NCYL=14,TYPE=N,DSF=N
DPD VOLID=DOSRES,CYL=479,NCYL=162,TYPE=N,DSF=N
DPD VOLID=SYSWK1,CYL=1500,TYPE=N,DSF=N
DLA NAME=AREA1,VOLID=DOSRES,CYL=71,NCYL=3,DSF=N
SVA SDL=600,GETVIS=(2MK,8M),PSIZE=(256K,4M)
```

Figure 12. VSE IPL Procedure

B.2 VSE Allocation Procedure

```
  CATALOG ALLOC.PROC   DATA=YES      REPLACE=YES
ALLOC BG=1536K
SIZE BG=1280K
ALLOC F1=1600K
SIZE F1=768K
ALLOC F2=10M
SIZE F2=3M
ALLOC F3=20M
SIZE F3=600K
ALLOC F4=75M
SIZE F4=8M
ALLOC F5=75M
SIZE F5=8M
ALLOC F6=512K
SIZE F6=256K
ALLOC F7=512K
SIZE F7=256K
ALLOC F8=75M
SIZE F8=4M
ALLOC F9=75M
SIZE F9=4M
ALLOC FA=512K
SIZE FA=256K
ALLOC FB=512K
SIZE FB=256K
SYSDEF DSPACE,DSIZE=20M
NPGR BG=100,F2=255,F3=100,F4=100,F5=50,F6=50,F7=50,F8=200
NPGR F9=50,FA=50,FB=50
/+
```

Figure 13. VSE ALLOC Procedure

B.3 SVA Load Job for LE/VSE

```
* $$ JOB JNM=SVALE,CLASS=6,DISP=D
* $$ LST CLASS=A,DISP=H
// JOB SVALE   CATALOG SDL PROCEDURE
// EXEC LIBR
ACCESS SUBLIB=IJSYSRS.SYSLIB
CATALOG SVALE.PROC  DATA=YES REPLACE=YES
// LIBDEF PHASE,SEARCH=(PRD2.SCEECICS,PRD2.SCEEBASE)
SET SDL
CEEPLPKD,SVA
CEEPLPKA,SVA
CEEQMATH,SVA
CEEYDTS,SVA
IGZERRE,SVA
CEEEV005,SVA
/*
/+
/*
/&
* $$ EOJ
```

Figure 14. SVA Load Job for LE/VSE

B.4 VTAM Startup Job

```
* $$ JOB JNM=VTAMSTRT,DISP=D,CLASS=3
// JOB VTAMSTRT  START VTAM
// OPTION DUMP,SADUMP=5
// SETPARM XNCPU=¢¢
// EXEC PROC=$COMVAR,XNCPU
// EXEC DTRSETP,PARM=¢CPUVAR&XNCPU;;SET XSTATF3=ACTIVE¢
/*
// SETPFIX LIMIT=424K
// ASSGN SYS000,UA
// ASSGN SYS001,DISK,VOL=SYSWK1,SHR     TRACE FILE ASSIGNMENT
// ASSGN SYS004,DISK,VOL=SYSWK1,SHR     TRACE FILE ASSIGNMENT
// ASSGN SYS005,DISK,VOL=SYSWK1,SHR     NCP LOAD/DIAG FILE ASSGN
// LIBDEF *,SEARCH=(PRD2.COMM,PRD2.COMM2,PRD2.CONFIG,                  *
               PRD1.BASED,PRD1.BASE),PERM
// LIBDEF   DUMP,CATALOG=SYSDUMP.F3,PERM
// EXEC ISTINCVT,SIZE=ISTINCVT,PARM=¢CUSTNO=0000-000-0000,VTAMPW=AXXN-N*
               XXE-XXXX-TXXI-NXXG¢,DSPACE=6M
// EXEC DTRSETP,PARM=¢CPUVAR&XNCPU;;SET XSTATF3=INACTIVE¢
/*
/&
* $$ EOJ
```

Figure 15. VTAM Startup Job

B.5 CICS Startup Job

```
* $$ JOB JNM=CICSV2,DISP=D,CLASS=4,EOJMSG=YES
* $$ LST CLASS=A,DISP=H,RBS=100
// JOB CICSV2  GNITENNA STARTUP OF SECOND CICS
// PAUSE ..........
// OPTION SADUMP=5
// LIBDEF *,SEARCH=(PRD2.RAMPCV,PRD2.SCEECICS,PRD2.SCEEBASE,            *
               PRD2.CONFIG,PRD1.BASED,PRD1.BASE,PRD2.RAMPC,             *
               PRD2.PROD,PRD2.CICSR,PRD2.DBASE),PERM
// LIBDEF DUMP,CATALOG=SYSDUMP.F4
// SETPARM TPMODE=¢¢
// SETPARM XMODEF4=¢¢
// SETPARM XNCPU=¢¢
// EXEC PROC=$COMVAR,XNCPU
// EXEC DTRSETP,PARM=¢CPUVAR&XNCPU;;SET XSTATF4=ACTIVE¢ **F4 ASSUMED
/*
// EXEC PROC=CPUVAR&XNCPU,TPMODE,XMODEF4                **F4 ASSUMED
// SETPFIX LIMIT=1024K
LOG
// ID USER=FORSEC,PWD=UIYKIB
NOLOG
// EXEC PROC=DTRCICS2                  LABELS FOR CICS FILES
// EXEC PROC=RAMPFILE                  LABELS FOR RAMP-C FILES
*   WAITING FOR VTAM TO COME UP
// EXEC IESWAITT
// IF XMODEF4 = COLD THEN                                **F4 ASSUMED
// GOTO COLD
// EXEC DFHSIP,SIZE=10M,PARM=¢SIT=CV,$END¢,DSPACE=4M
/*
// GOTO END
/. COLD
// EXEC DFHSIP,SIZE=10M,PARM=¢SIT=CV,START=COLD,$END¢,DSPACE=4M
* /*
/. END
// EXEC DTRSETP,PARM=¢CPUVAR&XNCPU;;SET XSTATF4=INACTIVE¢
/*
/&
* $$ EOJ
```

Figure 16. CICS Startup Job

B.6 VSAM Compressed File Definition

```
DEFINE CLUSTER ( -
       NAME (FILEA1C) -
       COMPRESSED -
       CYL(14      3) -
       SHAREOPTIONS (1) -
       RECORDSIZE (209    209) -
       NOIMBED -
       SPEED -
       VOLUMES (VSAM01) -
       NOREUSE -
       INDEXED -
       FREESPACE (15 7) -
       KEYS (4   0) -
       TO (99366)) -
       DATA (NAME (FILEA1C.@D@) -
       CONTROLINTERVALSIZE (4096) -
       INDEX (NAME (FILEA1C.@I@) -
              CISZ(2048)) -
       CATALOG (RAMPC.COMP.CATALOG)
```

Figure 17. VSAM Compressed File Definition

List of Abbreviations

AOR	Application Owning Region
AR	Attention Routine
CEC	Central Electronic Complex
CI	Control Interval
CMOS	Complementary Metal Oxide Semiconductor
CMT	CICS Maintained Data Table
CP	Central Processor
CPU	Central Processing Unit
CTC	Channel to Channel
DIM	Data In Memory
DSA	Dynamic Storage Area
ESDS	Entry Sequenced Data Set
FOR	File Owning Region
FS	Function Shipping
IBM	International Business Machines Corporation
II	Interactive Interface
IPL	Initial Program Load
I/O	Input/Output
ITSO	International Technical Support Organization
JA	Job Accounting
KSDS	Key Sequenced Data Set
LPAR	Logical Partition
LSR	Local Shared Resources
LU	Logical Unit
LUW	Logical Unit of Work
min	minute
Min	Minute
MRO	Multiregion Operation
NPA	Non-Parallel Application
NPS	Non-Parallel Share
NSR	Non-Shared Resources
PSW	Program Status Word
RAMP-C	Requirements Approach to Measuring Performance in COBOL
SD	Standard Dispatcher
SDAID	System Debugging Aid
SVA	Shared Virtual Area
TD	Turbo Dispatcher
TOR	Terminal Owning Region
TPNS	Teleprocessing Network Simulator
TR	Transaction Routing
UMT	User Maintained Data Table
VRDS	Variable-length Relative Record Data Set

Index

ITSO Technical Bulletin Evaluation — RED000

International Technical Support Organization
VSE/ESA Version 2 Release 1
The Turbo Dispatcher
February 1996

Publication No. SG24-4674-00

Your feedback is very important to help us maintain the quality of ITSO Bulletins. **Please fill out this questionnaire and return it using one of the following methods:**

- Mail it to the address on the back (postage paid in U.S. only)
- Give it to an IBM marketing representative for mailing
- Fax it to: Your International Access Code + 1 914 432 8246
- Send a note to REDBOOK@VNET.IBM.COM

Please rate on a scale of 1 to 5 the subjects below.
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction ____

Organization of the book	____	Grammar/punctuation/spelling	____
Accuracy of the information	____	Ease of reading and understanding	____
Relevance of the information	____	Ease of finding information	____
Completeness of the information	____	Level of technical detail	____
Value of illustrations	____	Print quality	____

Please answer the following questions:

a) If you are an employee of IBM or its subsidiaries:

Do you provide billable services for 20% or more of your time? Yes____ No____

Are you in a Services Organization? Yes____ No____

b) Are you working in the USA? Yes____ No____

c) Was the Bulletin published in time for your needs? Yes____ No____

d) Did this Bulletin meet your needs? Yes____ No____

If no, please explain:

____

____

What other topics would you like to see in this Bulletin?

____

____

What other Technical Bulletins would you like to see published?

____

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

______________________________ ______________________________
Name Address

______________________________ ______________________________
Company or Organization

______________________________ ______________________________
Phone No.

Fold and Tape **Please do not staple** Fold and Tape

PLACE
POSTAGE
STAMP
HERE

IBM International Technical Support Organization
Department 3222, Building 71032-02
POSTFACH 1380
71032 BOEBLINGEN
GERMANY

Fold and Tape **Please do not staple** Fold and Tape

Printed in U.S.A.

SG24-4674-00