International Technical Support Organization SG24-4514-00

A Comparison of System/390 Configurations - Parallel and Traditional

October 1995

IBM

International Technical Support Organization
Poughkeepsie Center

International Technical Support Organization

SG24-4514-00

A Comparison of System/390 Configurations - Parallel and Traditional

October 1995

Take Note!

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

First Edition (October 1995)

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. 541 Mail Station P099
522 South Road
Poughkeepsie, New York 12601-5400

Abstract

There is a clear consensus, from IBM, other vendors and industry consultants, that parallel is the future for large systems. The question is not whether parallel systems will happen, but when a user should move to them.

This document compares parallel sysplex to traditional System/390 MVS configuration options. It shows the relative advantages of each, in terms of availability, performance, cost and systems management effort. This will help users to select the configuration that most closely meets their needs.

The configurations discussed are a single processor, a single multiprocessor, two processors sharing DASD for backup purposes, two processors in a parallel sysplex, and five processors in a parallel sysplex. The document is intended to assist users to answer questions such as:

- Should I implement a single stand-alone processor, or multiple processors in a parallel sysplex?
- Should I upgrade to a larger processor, or should I add an additional processor with a parallel sysplex?
- If I have multiple processors, is it more efficient to run them with shared DASD (and no parallel sysplex), or in a parallel sysplex?
- Should I run a processor as a single system, or partitioned, either physically, or logically using PR/SM?

This document was written for information systems managers and planners. It is also relevant to IBM large systems marketing and technical personnel. Some knowledge of System/390 and parallel sysplex is assumed.

(44 pages)

Contents

Figures

Tables

Special Notices

This publication is intended to help System/390 users to select between the different System/390 configuration options. The information in this publication is not intended as the specification of any programming interfaces. See the PUBLICATIONS section of the IBM Programming Announcements for System/390 products 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.

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

CICS	CICS/ESA
ES/9000	ESCON
IBM	IMS
MVS/ESA	OS/2
PR/SM	RMF
SP	S/390
System/390	SystemView
VM/ESA	VTAM
3090	

The following terms are trademarks of other companies:

Windows is a trademark of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

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

C-bus is a trademark of Corollary, Inc.

Other trademarks are trademarks of their respective companies.

Preface

There is a clear consensus, from IBM, other vendors and industry consultants, that parallel is the future for large systems. The question is not whether parallel systems will happen, but when a user should move to them.

This document compares parallel sysplex to traditional System/390 configuration options. It shows the relative advantages of each, in terms of availability, performance, lowest cost and systems management effort. This will help users to select the configuration that most closely meets their needs.

The configurations discussed are a single processor, a single multiprocessor, two processors sharing DASD for backup purposes, two processors in a parallel sysplex, and five processors in a parallel sysplex. The document is intended to assist users to answer questions such as:

- Should I implement a single stand-alone IBM 9672 CMOS processor, or multiple processors in a parallel sysplex?
- Should I upgrade to a larger processor, or should I add an additional processor with a parallel sysplex?
- If I have multiple processors, is it more efficient to run them with shared DASD (and no parallel sysplex), or in a parallel sysplex?
- Should I run a processor as a single system, or partitioned, either physically, or logically using PR/SM?

This document is based on the current IBM S/390 processor range. The techniques in this document could be applied to other processors by making appropriate changes for product differences.

This document was written for information systems managers and planners. It is also relevant to IBM large systems marketing and technical personnel. Some knowledge of System/390 and parallel sysplex is assumed.

How This Document is Organized

The document is organized as follows:

- Chapter 1, “Introduction and Summary”

  This introduces the document, and summarizes its conclusions.

- Chapter 2, “Comparison of Various Configurations”

  This describes the System/390 configurations that we compared, the assumptions we used, and the rating method. The comparisons are then presented in Table 1 on page 7.

- Chapter 3, “Explanation of Availability Items”

  This explains the availability line items and ratings in Table 1 on page 7.

- Chapter 4, “Explanation Of Performance Items”

  This explains the performance line items and ratings in Table 1 on page 7.

- Chapter 5, “Explanation of Lowest Cost Items”

  This explains the cost line items and ratings in Table 1 on page 7.

- Chapter 6, "Explanation of System Management Effort Items"

  This explains the system management effort line items and ratings in Table 1 on page 7.

- Appendix A, "Calculation of Usable Capacity"

  This shows a sample calculation of the usable capacity for configurations of each type considered in Table 1 on page 7.

Related Publications

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

- *System/390 MVS Sysplex Overview: Introducing Data Sharing and Parallelism in a Sysplex*, GC28-1208
- *System/390 MVS Sysplex Systems Management*, GC28-1209
- *System/390 MVS Sysplex Hardware and Software Migration*, GC28-1210
- *System/390 MVS Sysplex Application Migration*, GC28-1211
- *MVS/ESA SP V5 Planning: Workload Management*, GC28-1493
- *CICS/ESA Dynamic Transaction Routing in a CICSplex*, SC33-1012
- *Large Systems Performance Reference*, SC28-1187

International Technical Support Organization Publications

- *System/390 MVS Parallel Sysplex Performance*, GG24-4356
- *MVS/ESA Sysplex Migration Guide*, GG24-3925

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.

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 Master Cards are accepted. Outside the USA, customers should contact their local IBM office. Guidance may be obtained by sending 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 (such as WebExplorer from the OS/2 3.0 Warp BonusPak) to the following:

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

IBM employees may access LIST3820s of redbooks as well. Point your web browser to the IBM Redbooks home page:

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

Other Related Material

The following documents are available to IBM employees:

- *Large Systems Performance Reference (LSPR)*, obtainable from MKTTOOLS locally, or by typing the following command from your local VM user ID:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LSPRPC PACKAGE
  ```

- *S/390 Parallel Sysplex System Performance Considerations Presentation Guide*, obtainable from MKTTOOLS locally, or by typing the following command from your local VM user ID:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET MFTS PACKAGE
  ```

- *CICS/ESA 4.1 Performance Comparison With CICS/ESA 3.3*, obtainable from MKTTOOLS locally, or by typing the following command from your local VM user ID:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET CICSE41P PACKAGE
  ```

Acknowledgments

This project was designed and managed by:

Brian Macfaden
International Technical Support Organization, Poughkeepsie Center

The authors of this document are:

Luigi Ingrosso
IBM Italy

Ramada Pereira
IBM Portugal

This publication is the result of a residency conducted at the International Technical Support Organization (ITSO), Poughkeepsie Center.

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

Ray Berard	IBM Large Scale Computing - Europe
Angelo Corridori	IBM S/390 Division
Joan Kelley	IBM S/390 Division
Howard Marin	IBM S/390 Division
Mary Moore	IBM S/390 Division
Vincent Mthimunye	IBM South Africa
Madeline Nick	IBM S/390 Division

Chuck Poland	ITSO, Poughkeepsie Center
Tom Rosamilia	IBM S/390 Division
Tom Russell	ITSO, Poughkeepsie Center
Lisa Spainhower	IBM Systems Technology and Architecture Division
Nora Sweet	IBM S/390 Division

Chapter 1. Introduction and Summary

The purpose of this document is discussed in the “Preface” on page xi. You may wish to review it if you have not already done so. This chapter introduces the document and summarizes its conclusions.

1.1 Introduction

In this document, we compare five types of System/390 configurations with approximately the same total capacity:

- Three using traditional non-parallel configurations:
  - A single image processor.
  - A single image multiprocessor that can be physically partitioned.
  - Two separate processors sharing DASD, so that the processors can back each other up. Only system data sets are shared between the processors.
- Two using parallel sysplex to provide service to a common group of users with parallelism and dynamic workload balancing:
  - Two processors
  - Five (smaller) processors

In order to provide a simple summary, we have taken a view of a *typical* System/390 user environment. Some of the comparisons may therefore be incorrect in a particular user′s environment. We have explained our thought process for each item. Therefore, you can use this document as a starting point, and modify the comparison for your situation where necessary.

The comparisons in this document are summarized in Table 1 on page 7.

1.2 Summary

A parallel sysplex environment offers many advantages over previous non-parallel environments. It can help you lower your total cost of computing while providing you with a level of availability previously unattainable. At the same time, advanced new functions provide good performance and easier management of multiple processors.

Both a single processor and a parallel sysplex provide better use of resources and lower cost than multiple independent processors. However, the availability of a parallel sysplex far exceeds the availability of any other configuration.

In addition, a parallel sysplex configuration can provide greater value to your business in the following areas:

- Data is a critical business resource. The ability to access and use that data, whenever it is needed, is therefore of great value. Parallel sysplex can provide continuous access to your data and applications.
- Change is a healthy sign for a business. Business growth and maintaining a competitive edge require changes to applications, and upgrades of software and hardware. It is important that these changes are made in a timely manner. Parallel sysplex allows changes to be made while critical

applications are still running, thus allowing a higher rate of change and more rapid implementation.

- Businesses are under cost pressure to only implement capacity as they require it. With a parallel sysplex, additional processors can be added as required. These processors can be of different sizes, and can therefore match the required capacity much more closely.

Our main findings are:

Availability

- Any application using parallel sysplex data sharing and dynamic workload balancing will continue to run if any system, subsystem or application fails, or is removed for hardware or software changes.
- The more systems there are in the parallel sysplex, the less the impact of a planned outage or failure of one of those systems.

Performance

- Parallel sysplex is likely to be more efficient than two processors sharing DASD. This is because a parallel sysplex is able to treat multiple systems as one set of resources.
- Parallel sysplex may also be more efficient than a single processor. Its coupling efficiency can be better than the multiprocessor efficiency of a single processor, especially as the number of engines increases.
- A single processor has a limited maximum capacity; a parallel sysplex can grow much larger.
- The response time of a parallel sysplex configuration is comparable to that of other configurations. Even if the parallel sysplex consists of systems with less processor engine power, the response time difference for most applications is insignificant.

Cost

- There is an initial investment cost for parallel sysplex configurations. This is quickly offset by growth flexibility - with additional upgrade options and the ability to install capacity as it is needed - and software cost savings.
- One system has a lower cost than two systems sharing DASD.

System Management Effort

- Parallel sysplex requires less effort to manage than multiple independent processors, as it provides many benefits of a single system image.
- A single processor takes less effort to manage than two processors.
- Parallel sysplex and a single system both provide benefits of a single system image. Each has some advantages over the other, with parallel sysplex gaining from the resilience of the configuration to continue running through planned outages, failures and performance issues. These reduce the effort to plan and manage such situations. Many systems management products, from both IBM and other vendors, support management of a parallel sysplex.

Chapter 2. Comparison of Various Configurations

This chapter describes the System/390 configurations that we compared, the assumptions we used, and the rating method. The comparisons are then presented in Table 1 on page 7.

IBM System/390 MVS users have a number of alternatives in the way they configure and run their systems. Parallel sysplex is one of those alternatives. A user can even have multiple options with the same processor hardware, for example:

- Two separate processors could have no shared DASD, they could share a common DASD subsystem, or they could be connected together in a parallel sysplex.
- An IBM ES/9000 9021-962 6-way multiprocessor could be configured as a single processor, or physically partitioned into two 9021-831 type processors that could share DASD or could be in a parallel sysplex.

To help users decide how to run their S/390 processors, and to help them select the configuration they wish to buy, we have compared several configuration options.

2.1 Terminology

In this document, we use the term *shared DASD* for multiple systems able to share a DASD subsystem at a *data set* level, using GRS (or reserve/release). The objective of this level of sharing is to share some system data sets, such as the security database and the JES2 spool, and to provide a capability for the systems to back each other up. We assume that databases are not being shared. An example of this is shown in the left hand diagram of Figure 1. Production work runs on one processor, and development work runs on the other. There is minimal sharing of the DASD subsystem. However, if the production processor fails, the production work can be restarted on the development processor, as it is connected to the production DASD.

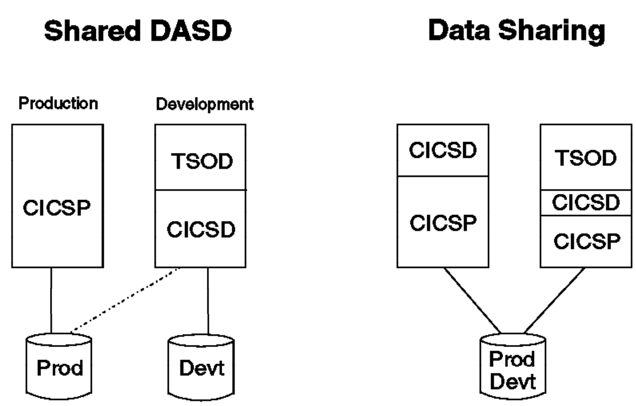

Figure 1. Shared DASD Compared to Data Sharing

We use the terms *parallel sysplex* and *data sharing* for multiple systems sharing a DASD subsystem and databases at a *record* level, using a coupling facility to provide record locking between the systems. An example of this is shown in the right hand diagram of Figure 1. Production and development CICS systems can be split across the two processors, accessing a common database. This provide greater resilience to a system failure.

2.2 The Configurations

Apart from the last column (*five systems parallel*) in Table 1 on page 7, all of the configurations are at the same technology level. They could be all IBM 9021-711 based processors, or all 9672 Model R2 processors, for example. These columns can be used to compare configuration differences. We do not compare, for example, a 3090-J model to a 9672. For your particular situation, you may also need to take account of differences in processor technologies (such as acquisition cost, maintenance cost, running costs, and software licensing options).

The column *five systems parallel* can be used with the rest of the table to compare CMOS processors in a parallel sysplex to bipolar processor configurations. This is included as it is likely to be of interest to many users. It shows the effect of having more, smaller processors compared to fewer larger ones.

These configurations are not intended to be a complete set. For example, a parallel sysplex can have a mixture of processor technologies, with bipolar and CMOS processors sharing data. You should be able to evaluate other configurations from the ones we have analyzed. Similarly, while we give examples of IBM processors for each configuration, you could use this methodology to evaluate other types of processors by taking account of product differences.

We assume that each configuration has a similar total capacity.

The configurations we discuss in Table 1 on page 7 are:

One system: This is a single processor that cannot be physically partitioned. An example would be an IBM 9672-R42 (four-way CMOS processor). For this discussion we assume that the processor runs a single MVS production system. It may also run some test systems, with PR/SM or VM/ESA providing the management of the different systems. The key points about this configuration are:

- There is only one production MVS system to manage. Any sharing of data between production and test MVS systems is minimal.
- All processor resources are under the control of one resource manager - MVS, VM/ESA or PR/SM.

If you run multiple production MVS systems on a single processor, your configuration is more like that in *two systems*.

One mp: Read as *one multiprocessor*. This is a single processor that can be physically partitioned into two separate processors if required. An example would be an IBM ES/9000 9021-962 (6-way bipolar multiprocessor), that can be physically partitioned into the equivalent of two 9021-831s (3-way processors). For this discussion we assume that the processor runs a single MVS production system. It may also run some test systems, with PR/SM or VM/ESA providing the management of the different systems. The key points about this configuration are:

- The processor normally runs in single image mode.
- There is only one production MVS system to manage. Any sharing of data between production and test MVS systems is minimal.
- All processor resources are under the control of one resource manager - MVS, VM/ESA or PR/SM.
- The processor is configured so that if one side fails, it can be physically partitioned to allow production work to run on one side while the other side is available for repair. We therefore assume that all DASD, and other devices, are connected to both sides of the processor.

In practice, users of a very large processor often run multiple production MVS systems on it. They do this to improve manageability or availability. In this case, the configuration is more like that in *two systems*.

Two systems: This is two processors that share DASD at a data set level, and are not connected in a sysplex. The shared DASD is used only for sharing some system data sets, and for backup if one processor should fail. There is no database record level sharing between systems. These could be two separate processors, or a multiprocessor, such as an IBM 9021-962, that is run in physically partitioned mode. We assume for simplicity that the two processors have equal capacity, and that there are two production MVS systems, one on each processor. If one processor fails, workload can be manually moved to the other processor.

Two systems parallel: Read as *a two system parallel sysplex*. This is two processors that share DASD and are connected in a parallel sysplex using an external coupling facility. For simplicity, we assume the two processors have the same capacity. Examples would be two IBM 9672-R22s (two-way CMOS processors), or an IBM 9672-E02 (2 processors), or two IBM ES/9000 9021-831s (3-way bipolar processors).

Five systems parallel: Read as *a five system parallel sysplex*. This is five CMOS processors that share DASD and are connected in a parallel sysplex using an external coupling facility. For simplicity, we assume the five processors have the same capacity. We assume each processor has a smaller capacity than in *two systems parallel*, so we need more of them to provide the same total capacity. Examples would be five IBM 9672 R32s (three-way processor), or an IBM 9672-E05 (5 processors).

We took the single system (configuration *one system*) as the base for the comparisons.

2.3 Other Assumptions

To provide a comparison of *processor* configurations, we made assumptions about the rest of the configuration:

- The I/O configuration is adequate for each processor configuration. The I/O subsystem is not constraining workload throughput.
- All I/O devices can be accessed from all processors. For configuration *five systems parallel* this requires either ESCON connections to DASD, or connection with fewer than four channel paths to each processor.

- The software products and levels are the same in all configurations, and are capable of data sharing and dynamic workload balancing. We do not consider differences due to different software levels.
- In the parallel sysplex configurations, 50% of the workload uses data sharing and dynamic workload balancing, and is capable of being dynamically routed across the parallel sysplex. The rest of the workload is always associated with a particular processor.

2.4 The Rating Method

We compare the characteristics of each configuration in these areas:

Availability The effect of failures and system changes on the users′ ability to use applications.

Performance The capacity available to run applications, online response times and batch runtimes.

Lowest Cost The hardware and software costs.

System Management Effort The cost of people to manage the systems.

The items we consider to be of higher importance in each area are marked with a ♦ in column one of the table.

We rate each configuration using the following symbols:

0	is the base value. We took *one system* as the base.
+	is better than 0
+ +	is better than +
-	is not as good as 0
- -	is not as good as -

In rating each item, we consider ongoing costs and effort. Migration costs are excluded because they vary depending on your starting point.

This method of rating is somewhat simplistic, but it was chosen so that the differences between the configurations are easily seen in the table.

The ratings are the opinions of the authors. In the rest of this document, we explain how we rated each line in the table. The third column in the table shows the page number where that line is discussed. You can therefore follow our thought processes, and amend the table as necessary for your particular circumstances.

2.5 The Comparison Table

Table 1. Comparison of the Configurations							
Configuration (mp = multiprocessor)		**Page**	**One system**	**One mp**	**Two systems**	**Two systems parallel**	**Five systems parallel**
Technology type			Same technology				Different
AVAILABILITY							
♦	Scheduled Processor Upgrade	10	0	0	+	+ +	+ +
	Scheduled DASD Upgrade	11	0	0	+	+ +	+ +
♦	Scheduled Software Change	11	0	0	+	+ +	+ +
♦	Scheduled Environment Change	12	0	0	+	+ +	+ +
	Failure of Processor Hardware	13	0	+	+	+ +	+ +
	Failure of I/O Hardware	13	0	0	0	0	0
	Failure of Operating System	14	0	0	+	+ +	+ +
	Failure of Subsystems	14	0	0	0	+ +	+ +
♦	Failure of Application	15	0	0	0	+ +	+ +
♦	Failure Due To Operational Error	15	0	0	0	+	+
♦	Failure Of Environment	16	0	0	+	+ +	+ +
PERFORMANCE							
♦	Effective Use of Capacity	17	0	0	-	0	0
	Consistent Online Response Times	18	0	0	-	0	0
	Batch Job Runtimes	19	0	0	-	0	-
LOWEST COST							
♦	Processor Upgrade	21	0	0	0	+ +	+ +
	Coupling Hardware	22	0	0	0	-	- -
	DASD	23	0	0	-	-	-
	Consoles	23	0	-	-	0	0
	I/O Connectivity	24	0	-	-	-	- -
♦	Hardware Running Costs	24	0	0	0	0	+ +
♦	IBM Software License Charges	25	0	0	-	+	+ +
SYSTEM MANAGEMENT EFFORT							
	Operational Complexity	27	0	-	- -	-	-
	Automated Operations	27	0	0	- -	-	-
♦	Change Management	28	0	0	0	+ +	+ +
	Problem Management	29	0	0	-	0	0
	Workload Management	29	0	0	- -	-	-
	Capacity Planning	30	0	0	- -	+ +	+ +
	Performance Management	31	0	0	-	+ +	+ +
	Configuration Management	31	0	-	-	-	- -

Ratings:	**- -**	**-**	**0**	**+**	**+ +**
	Least good	Less good	Base	Better	Best

(*A Comparison of System/390 Configurations - Parallel and Traditional*, SG24-4514-00).

Chapter 3. Explanation of Availability Items

This section explains the reasoning behind the evaluation of the availability lines in the table.

We consider that the system is available if the processor, operating system, subsystems, applications and data that a user needs are available when required. We rate the different configurations on their likely availability as seen by the users of the systems. To do that, we look at the common causes of unavailability, and evaluate their likelihood of occurring, the likely duration of an outage, and the impact on the users.

Customer studies done by IBM suggest that the primary causes of unavailability are those marked with a ♦ in column one of Table 1 on page 7. You may wish to analyze your own causes of unavailability, to understand their relative importance in your environment,

We do not discuss outages for database reorganization. In all configurations, database reorganization affects those applications that use that database.

We will first describe the availability characteristics of a parallel sysplex, as these may not be familiar to you. Then we discuss scheduled outages, needed to make changes to the systems, and then unscheduled outages, generally caused by failures. In many installations, scheduled outages cause more loss of availability than unscheduled outages.

3.1 Parallel Sysplex and Availability

From the end users′ point of view, the parallel sysplex can provide continuous availability.

Parallel sysplex uses coupling technology to allow multiple processors to work in a single system image. A parallel sysplex can be configured for various levels of availability. For this study we assume a configuration that has full redundancy, and therefore no single points of failure. Failure of one component of the parallel sysplex can therefore be transparent to applications and users.

To explain how this works, let us assume a catastrophic failure of one of the processors in a parallel sysplex. One of the functions of a parallel sysplex is dynamic workload balancing. CICS transactions can be dynamically routed across multiple MVS images. If one of the systems fails, or suffers degraded performance, that is detected, and subsequent work is automatically routed to the other systems. A user of an application using dynamic workload balancing will only see an outage if either:

1. They had a transaction running on the failing system at the time of failure. In this case, they will need to re-enter the transaction.
2. They were logged on to a terminal region (for example a CICS TOR) on the system that failed. In this case, they will need to log on again. They can immediately log on to the same application name, using another feature of parallel sysplex; VTAM generic resources allows a user to log on to a generic application name. VTAM will connect the user to one of several applications on different processors, and will select one that is available at that time.

Thus we see that users of applications using dynamic workload balancing will see at most a short outage. The workload can be automatically picked up on the remaining systems.

Applications not using dynamic workload balancing will see an outage, and users will need to log on again. However, if the application uses data sharing, it can be restarted immediately on another system without waiting for the failing system to be recovered, as lock information is held in the coupling facility.

Automatic Restart Manager (ARM) is an MVS component that can restart a failing subsystem automatically, either on the same system or on another system. This can speed up the recovery of failing applications.

We will now look at the impact of different types of outage on the various configurations.

3.2 Scheduled Processor Upgrade

This item includes all processor hardware upgrades or changes that require a processor to be taken offline. This may include increasing the processor capacity, adding new features, or upgrading microcode. Note that the outage may be a number of hours long. Many users find it difficult to schedule hardware upgrades because of the outage duration.

With only one processor (configuration *one system*), no systems are available for the duration of the outage.

With a single partitionable system (configuration *one mp*), no systems are available for the period of the outage when both sides are required for the upgrade. If part of the upgrade can be done one side at a time, then the system can be partitioned. This may involve one or more restarts of the system, when all applications are unavailable, and at most 50% of the capacity would be available for the duration of the upgrade.

With two independent processors (configuration *two systems*), each system can be upgraded independently, leaving one system available at all times. Users will be disrupted while work is moved from the system to be upgraded, and back at the end of the upgrade. Only 50% of the capacity would be available for the duration of the upgrade.

With two processors in a parallel sysplex (configuration *two systems parallel*), many users will not see any outage. Applications using dynamic workload balancing will automatically route transactions to the system remaining. Users logged on to a terminal region on the system being upgraded will need to log off and log on again, when they will automatically be connected to a terminal region on the remaining system. If those applications use data sharing, the users can move across to the other system at their convenience, since copies of the application could be running on both systems at the same time. This outage will be much shorter than in *two systems*, as they do not need to wait for all users to log off and then the complete application to be moved. Only 50% of the capacity would be available for the duration of the upgrade. When the upgraded system is brought back into the parallel sysplex, transactions can be routed to it, and the capacity used, without users needing to log off and log on again.

Five processors in a parallel sysplex (configuration *five systems parallel*) is exactly the same as for configuration *two systems parallel*, except that 80% of the capacity remains available. Because only 20% of the capacity is affected, this provides more flexibility in scheduling the upgrade. In fact, the more images there are in a parallel sysplex, the less the impact of an outage.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	+	+ +	+ +

3.3 Scheduled DASD Upgrade

This includes changes or upgrades to the DASD subsystem that require operating systems or subsystems to be taken offline. This may be to move system volumes to a new string of DASD, for example. DASD changes are infrequent, but are often long.

Some DASD changes are the same for all configurations. Changes that involve only application data will affect those applications for all configurations. If the changes are nondisruptive, for example using IBM 3990 Dual Copy, then none of the systems are affected for all configurations. Such changes are therefore not considered.

The changes we are considering involve taking down a system or subsystem, and are therefore the same as 3.4, "Scheduled Software Change."

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	+	+ +	+ +

3.4 Scheduled Software Change

This includes all software maintenance activities - implementing a new software release or making a minor change - to the operating system, subsystems or an application. We assume the change has been tested, and we now need to put it into production. With only one MVS system (configurations *one system* and *one mp*), the affected system must be stopped and restarted.

With two independent MVS systems (configuration *two systems*), each system can be upgraded independently, but both systems must be upgraded. The effect on the users is the same as for one system, unless a long outage is required. In that case the ability to always have one system available may reduce the outage, though workloads will then need to be moved between the two systems at the start and end of the outage.

With two MVS systems in a parallel sysplex (configuration *two systems parallel*), many changes can be made by introducing them onto one of the systems first. Applications using dynamic workload balancing can continue to run while one system is removed, upgraded, and returned to the parallel sysplex. The new and old software versions will be able to run together in the parallel sysplex. If there is a problem with the new version, affected transactions can be automatically routed to the old version, thus masking the problem from the users. Then, when the new version is fully proven, the other system can be upgraded nondisruptively in the same way.

Five processors in a parallel sysplex (configuration *five systems parallel*) is exactly the same as for two processors in a parallel sysplex, except that only 20% of the capacity is affected at any time. This provides more flexibility in scheduling the upgrade. For software changes, which may be frequent, this is a significant advantage.

For these reasons we have rated:

One system	**One mp**	**Two systems**	**Two systems parallel**	**Five systems parallel**
0	0	+	+ +	+ +

3.5 Scheduled Environment Change

This is an upgrade or change to the computing center environment, such as power, air conditioning and chilled water. These often account for surprisingly large amounts of unavailability, as the outages tend to be long.

All of the configurations are affected the same by this type of outage, except that for CMOS processors neither air conditioning nor chilled water are needed. The IBM CMOS 9672 processors can have a Local Uninterruptible Power System (Local UPS), IBM machine type 9910, that can provide full power for a few minutes, to avoid an IPL of the system for a short power outage. Also, a UPS for CMOS configurations can be much less powerful, and therefore much cheaper, than for bipolar configurations, so may be more likely to be implemented. Note that to provide uninterrupted service, the I/O configuration must also be protected by a UPS.

Multiple processors (configurations *two systems*, *two systems parallel*, and *five systems parallel*) may allow more flexibility in configuring for environmental outages, by placing the processors in different computer centers. Then an outage at one computer center will only affect a subset of the configuration. Parallel sysplex will allow a nondisruptive move of dynamic workload balancing workloads between computer centers, as in 3.2, “Scheduled Processor Upgrade” on page 10.

For these reasons we have rated:

One system	**One mp**	**Two systems**	**Two systems parallel**	**Five systems parallel**
0	0	+	+ +	+ +

3.6 Failure of Processor Hardware

The failure of a processor is an infrequent event with today′s hardware. IBM′s latest bipolar processors are extremely reliable, and the IBM 9672 CMOS processors are expected to have even better reliability. However, if a processor does fail, it can take a significant amount of time to repair it. With only one processor (configuration *one system*), all applications are unavailable.

With a single partitionable system (configuration *one mp*), the processor can be partitioned so that workload can be run on one side while the other side is repaired. In our example, only 50% of the capacity would be available for that time.

With two independent processors (configuration *two systems*), the failure will affect one of them, and all applications on that processor will fail. The user must then decide which applications to continue running on the available system, and restart any of those that failed. Users of those systems will need to log on again after the applications are restarted. Only 50% of the capacity would be available until the failing processor is repaired, and then there would be a second disruption to some applications to move them back to the repaired processor.

With two processors in a parallel sysplex (configuration *two systems parallel*), all subsystems on the failing processor will fail. However, users of applications using dynamic workload balancing will see either no outage or only a very short outage - see 3.1, “Parallel Sysplex and Availability” on page 9. Other workload not capable of dynamic workload balancing can be moved to the running system as in configuration *two systems*. Only 50% of the capacity would be available until the failing processor is repaired. The system can automatically prioritize the workload on the capacity available.

Five processors in a parallel sysplex (configuration *five systems parallel*) is exactly the same as for two processors in a parallel sysplex, except that 80% of the capacity remains available.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	+	+	+ +	+ +

3.7 Failure of I/O Hardware

Failure of a non-DASD I/O device usually impacts all the configurations in the same way. For example, a printer failure reduces the printing throughput.

If configured with adequate redundancy, DASD failures have the same effect for all configurations. For example, we might configure two system residence volumes for all configurations, so that if one fails the other is still accessible. For the single system configurations (configurations *one system* and *one mp*) we would have one volume as a dual copy of the other, so that if either volume fails, there is no outage. For configuration *two systems* we would run both systems from the same system residence volume, and again have it dual copied. For the parallel sysplex configurations (*two systems parallel* and *five systems parallel*) we could again use dual copy, or we could run some systems off one system residence volume and some off the other one. Then, if one volume fails, systems using the other one will remain operational, and applications using dynamic workload balancing will see at worst a small outage.

Parallel sysplex introduces some additional external devices, namely a coupling facility and sysplex timer. To prevent service interruptions, configure two coupling facilities and two sysplex timers. If a sysplex timer fails, the second one will be used. If a coupling facility fails, the systems in the parallel sysplex will rebuild the structures in the backup coupling facility.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	0	0	0

3.8 Failure of Operating System

We include in this item all failures that effectively disable a whole MVS system, including failures of MVS, JES and VTAM. Such failures tend to be infrequent. An operating system failure has the same effect as a processor failure (see 3.6, "Failure of Processor Hardware" on page 13), except that the recovery action is to reIPL, which takes less time than a hardware repair. Applications will probably be restarted on the same system, and not moved to another system.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	+	+ +	+ +

3.9 Failure of Subsystems

Examples of subsystems are CICS and IMS. In all cases we assume that the subsystem is restarted on the processor it was running on, if possible, since that avoids the need to reschedule other workloads. As discussed in 3.1, "Parallel Sysplex and Availability" on page 9, the parallel sysplex configurations will provide no, or minimal, service break to the system users.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	0	+ +	+ +

3.10 Failure of Application

Some application errors will cause the whole application to fail in all configurations. However, many application errors only occur when a special set of circumstances comes together, or when some other software is upgraded. The application may run perfectly for many years, and then suddenly fail. In the non-parallel sysplex configurations, the application is not available until it is restarted.

In a parallel sysplex, for an application using dynamic workload balancing, the transaction routing program (for example, CICSPlex System Manager/ESA) can detect that a transaction failed on a particular system, so the next time the transaction is submitted CICSPlex System Manager/ESA can try it on a different system. The transaction is likely to work on the different system because the circumstances are different, or because that system is still at a previous software level. Thus, the transaction remains available to users while the problem is being fixed.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	0	+ +	+ +

3.11 Failure Due To Operational Error

This includes all possible operational errors, such as a command wrongly submitted by a system or subsystem operator, or a job submitted incorrectly by an operator or by an automated job scheduling package.

On a single system (configurations *one system* and *one mp*), an operator error can cause the whole system to fail.

As the complexity of the configuration increases, the operational tasks become more difficult. With two systems (configuration *two systems*) an error may only affect one of the systems, but the complexity may cause additional errors, or slower recovery. We therefore rate this the same as a single system.

In a parallel sysplex configuration, the operators benefit from:

- A single operator interface to all systems, for many products
- A single Hardware Management Console (IBM 9672)
- Continuous operations, where applications can continue to run even when an incorrect action impacts one of the processors

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	0	+	+

3.12 Failure Of Environment

This is a failure of the computing center environment, such as power, air conditioning and chilled water. These often account for surprisingly large amounts of unavailability, as the outages tend to be long. The considerations are the same as for 3.5, "Scheduled Environment Change" on page 12, and the ratings are the same.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	+	+ +	+ +

Chapter 4. Explanation Of Performance Items

This section explains the reasoning behind the evaluation of the performance lines in the table. We discuss the usable capacity of the system, the response time for online transactions, and the runtime for batch jobs.

4.1 Effective Use of Capacity

One of the key measures of a processor is its capacity, that is, the maximum amount of work that can be processed through the system. This is measured in number of online transactions per second, or number of completed batch jobs per hour.

IBM has measured the capacity of many processors, and has measured the effect of data sharing. This item in the table uses that information to consider the amount of work a user could realistically process through each configuration.

Whenever multiple processor engines are used to process a common workload, some of the processor capacity is used for intersystem communication, to keep the processor engines synchronized:

- In a single system (configurations *one system* and *one mp*), this is known as the multiprocessor effect, and it shows in IBM's processor measurements. The key factor is the multiprocessor ratio for the processor. This is the ratio of the capacity of the processor to the number of processor engines multiplied by the capacity of a single engine processor of the same type.
- With two separate systems (configuration *two systems*), the intersystem communication is often through Global Resource Serialization (GRS). However, there is no ability to dynamically move workload between these two separate systems, which means that the available capacity may not be available to the workload that needs it.

  As an example to help explain this concept, an IBM 9021-962 can be partitioned into two physical 9021-831s. While two 831s have a greater aggregate capacity than a 962, because of the multiprocessor effect, for most workloads the 962 will give better throughput. On the 962, one version of MVS can optimize the use of all six processor engines across the total workload.

  The key factor is how efficiently the workload can be manually balanced, and with how much effort.
- In a parallel sysplex, the intersystem communication uses the coupling facility. The amount of communication depends on the amount of data being shared between the processors in the parallel sysplex. However, dynamic workload balancing with data sharing allows the workload to be scheduled across the whole parallel sysplex, and thus removes the workload balancing issues of two separate systems. In addition, CICS intersystem communication (ISC) connections can be converted to CICS multiregion operation (MRO) connections, with a significant reduction in processing capacity used.

  The key factor is the amount of data sharing activity.

To help in making this assessment, we calculated the usable capacity for samples of each configuration. These calculations are in Appendix A, "Calculation of Usable Capacity" on page 33.

One system is rated better than two systems sharing DASD, as the single MVS system will almost always be more efficient than manual workload placement across two MVS systems.

When comparing one system to two systems in a parallel sysplex, our sample calculations show that either can be better, depending on the multiprocessor and data sharing efficiencies of the processor type (see A.9, "Conclusions" on page 40). We therefore rate these the same. However, note that the capacity of a single system is limited to 10-12 processor engines, while a parallel sysplex can have a capacity of 32 MVS images.

Comparing five CMOS systems to two bipolar systems, the more efficient data sharing outweighs the increased number of systems, so we rate it the same.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	-	0	0

4.2 Consistent Online Response Times

This evaluates the consistency of response times for online transactions. All of the configurations will provide similar response times when there are no resource constraints. We therefore look at the consistency of response times, for example at peak workload times of the day.

System/390 MVS Parallel Sysplex Performance, GG24-4356-00, contains response time measurements in a parallel sysplex. These are on page 54 (multiple 9672 processors), page 62 (9021 with multiple 9672 processors), and page 75 (multiple 9021 processors). From these measurements we make these observations:

- Transactions running on a slower processor (as in configuration *five systems parallel*) will have a longer response time than those running on a faster processor. The values measured were 0.2 seconds and 0.1 seconds CICS internal response time, respectively. Once network delay time is added, this difference will not be noticeable to users. Only for transactions using a lot of processor time may the difference be noticeable.
- If there is only one system, all resources are available to all transactions. If there are two separate systems, not in a parallel sysplex, transactions only have access to half of the resources, and are therefore more likely to be delayed by workload peaks.
- In a parallel sysplex, transactions are routed to the system according to their current performance, so response times will be better, and more consistent, than when running on two separate systems.
- With CICS for MVS/ESA Version 4, CICS multiregion operation (MRO) can use sysplex connections rather than VTAM connections. This can provide a performance improvement for the parallel sysplex configurations.

- With multiple systems, there will more DASD contention than with a single system, but for online workloads we do not expect this to be significant.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	-	0	0

4.3 Batch Job Runtimes

This evaluates the runtimes for batch jobs. We considered the following:

- If there is only one system, all resources are available to all jobs, and I/O contention is minimized. With two systems not in a parallel sysplex, there is greater I/O contention at the volume and controller level, and lockouts at the data set level. To minimize these effects, job scheduling restrictions may be imposed, thus lengthening the total batch runtime.
- With multiple systems in a parallel sysplex, databases can be shared between jobs running on different MVS systems, thus allowing additional scheduling flexibility. There is still likely to be more I/O contention than with a single system.
- Jobs running on a smaller processor (configuration *five systems parallel*) will have a longer processor time component than those running on a faster processor. However, for most batch jobs, the processor time is a small proportion of the total job runtime, often about 5%. Therefore, for most batch jobs, the runtime difference is negligible.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	-	0	-

Chapter 5. Explanation of Lowest Cost Items

This section explains the reasoning behind the evaluation of the lowest cost lines in the table. It includes the major hardware and software costs associated with a S/390 configuration. People costs are covered in Chapter 6, "Explanation of System Management Effort Items" on page 27.

We considered the hardware and software requirements for each configuration, and then rated the configurations in terms of their likely costs. Prices of hardware, in particular, continually change, so we did not attempt to calculate an actual cost for each configuration. You can use the discussion here to help you do that for your particular situation.

We rated these items with pluses for better and minuses for less good, to be consistent with the other items. Therefore pluses mean a lower cost, and minuses mean a higher cost.

5.1 Processor Upgrade

When a non-parallel sysplex configuration needs additional capacity, the simplest option is to upgrade one of the processors. While it is possible to add capacity by installing an additional processor, this requires a workload split, which is equivalent to moving from configuration *one system* to configuration *two systems*, with its associated costs.

Upgrading a processor will either be an upgrade of a current processor, which may be old technology, or a total replacement. Upgrading a current processor uses the technology of that processor. As costs reduce with newer technology, an older technology upgrade may be more expensive than other alternatives. Replacing a current processor may have cost implications because of current leases on that processor, or because of the book value of that processor in your accounting system.

A parallel sysplex configuration provides all of the options of a non-parallel sysplex configuration, plus another one. With multiple processors in a parallel sysplex, it is very easy to add an additional processor. Workloads using dynamic workload balancing can make use of the additional processor with very little effort, and management of an additional system is easy with the parallel sysplex single system image.

The new processor can be any coupling capable processor. This gives many options, in terms of both technology level and capacity. The ability to use the latest technology typically provides the best price/performance. The ability to use smaller processors improves your upgrade granularity. Capacity can be added in smaller increments, as needed, which minimizes both hardware and software costs. This is practical in a parallel sysplex because the capacity can be added without disrupting current applications. Leases and book value on current processors are also not affected. Thus parallel sysplex can provide the lowest upgrade cost.

In Figure 2 on page 37 we show how configuration *two systems* will require one processor to be upgraded even though the second processor is not at maximum capacity. A user with this environment will therefore need an upgrade earlier than the total available capacity may suggest. This is a major upgrade cost factor, but as it was included in 4.1, “Effective Use of Capacity” on page 17, we have not included it again here.

Backup systems may also involve additional costs. It may be necessary to upgrade both sides of configuration *one mp*, or to upgrade both processors in configurations *two systems* and *two systems parallel*, to provide sufficient capacity if one processor should fail. With many processors in the parallel sysplex, such as in configuration *five systems parallel*, the failure of any one processor only reduces the total system capacity by a small percentage, so additional resources for backup will probably not be required.

With a greater number of processors, the total amount of processor storage and the total number of channels will be greater, since each processor needs sufficient of these resources to handle its own peak requirement. However, these are likely to be cheaper on the most recent technology processors.

For these reasons we have rated:

One system	**One mp**	**Two systems**	**Two systems parallel**	**Five systems parallel**
0	0	0	+ +	+ +

5.2 Coupling Hardware

Parallel sysplex environments use additional hardware to enable parallelism and high speed data sharing. The initial investment cost for these items is quickly offset by other cost savings. Non-parallel sysplex environments do not provide these facilities and so do not require the extra hardware.

The extra hardware required is:

Sysplex timers Sysplex timers are required to synchronize the time of day (TOD) clocks between multiple independent processors. We assume two sysplex timers for each parallel sysplex configuration, to provide redundancy. If less redundancy is required, then one sysplex timer is sufficient.

Coupling facilities The number of coupling facilities required depends on the workload and on the need for connectivity. We would configure the same number of coupling facilities in all the parallel sysplex configurations, as the workload is the same. We assume a minimum of two coupling facilities in each case, to provide redundancy. If less redundancy is required, then one coupling facility may be sufficient.

Coupling links Coupling links connect each processor in a parallel sysplex to the coupling facilities. For availability, we would connect two links from each processor to each coupling facility. The number of coupling links required also depends on the workload, which is the same in each case. Since configuration *five systems parallel* has more processors than configuration *two systems parallel*, it will require more links.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	0	-	- -

5.3 DASD

This item discusses the DASD configuration differences between the options.

We assumed in 3.7, "Failure of I/O Hardware" on page 13 that all systems were configured for equivalent availability. For example, we assume that all configurations have two system residence volumes. On this basis, all configurations would have the same number of DASD volumes. However, some data sets, such as page data sets, cannot be shared between systems, so there would be a slight increase in DASD space required as the number of processors increases.

A parallel sysplex will generate some additional I/O activity. If one system updates a database record in its memory, other interested systems are notified. Next time they wish to access that record, they must read it from DASD or the coupling facility, rather than using the copy in their memory, which is now out of date. This may require a slight increase in some resources, such as DASD controller cache.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	-	-	-

5.4 Consoles

This item discusses the console differences between the options.

Non-parallel sysplex systems require one hardware console and at least one operating system console for each processor. Configuration *one mp* would require these consoles for each side, to allow it to run in partitioned mode.

In a parallel sysplex, one operating system console can operate all of the systems, thus reducing the total number of consoles required.

With IBM 9672 CMOS processors, a single hardware management console (HMC) can manage all of the processors.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	-	-	0	0

5.5 I/O Connectivity

This item discusses the I/O connectivity differences between the options. It covers connection of all the processors to I/O devices, and channel connections between processors. The costs are primarily in control unit channel adapters and in ESCON directors (channel costs are included in 5.1, "Processor Upgrade" on page 21). We assume, for maximum flexibility, that all processors are attached to all devices, wherever possible. We also assume ESCON attachment.

The single system configuration *one system* is the simplest. There is only one processor to connect to, and no communications between processors.

Configuration *one mp* only has a single processor. However, because it has two sides that can run independently, we assume that the I/O configuration is designed so that both sides are attached to all devices. This allows either side to run the workload, and effectively makes the I/O configuration very similar to that of configuration *two systems*.

Configuration *two systems* requires connections from two processors to all devices. This could be achieved either directly to channel adapters on the I/O control units, or through ESCON directors. Intersystem communication would be needed, to provide locking functions to control queueing on resources. Global Resource Serialization (GRS) provides this function. It would use channel to channel (CTC) connections between the processors.

Configuration *two systems parallel* is the same as configuration *two systems*, except that some intersystem communication can be routed through the coupling facility instead, which may reduce the number of CTC channels needed.

In configuration *five systems parallel*, five systems need to be connected to all I/O devices. Control units are generally not capable of having sufficient channel adapters to support five systems (though some DASD control units could if they only need three paths to each system), so connections would probably use ESCON directors. The number of switching ports and switching devices needed depends on the number of control units (the same in each configuration) and the number of processors, which is greater in this configuration.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	-	-	-	- -

5.6 Hardware Running Costs

This item includes the running costs associated with each configuration. These are:

- Hardware maintenance
- Electric power
- Uninterrupted power supplies
- Cooling equipment
- Machine room floor space, and need for raised floor

Parallel sysplex has minimal impact on these items. The only additions in a parallel sysplex are the coupling facilities and sysplex timers.

The technology of the processors does have a major impact. The configuration *five systems parallel* uses CMOS technology, and therefore has much lower requirements than the bipolar configurations in all of these areas.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	0	0	+ +

5.7 IBM Software License Charges

This item covers IBM software license charges.

With the introduction of parallel sysplex in April 1994, IBM introduced a new structure for software pricing, called Parallel Sysplex License Charge (PSLC). In a parallel sysplex configuration, software is priced according to the total processing capacity that it runs on, rather than the number of processors. This new pricing structure has two benefits for users of parallel sysplex:

- A lower overall software cost when compared to non-parallel sysplex environments
- The potential to further lower software costs by running some software products on a subset of the processors in the parallel sysplex

In September 1994, PSLC pricing was extended to single processor environments that satisfy certain hardware and software requirements that allow them to support a parallel sysplex.

In our ratings we assume that configurations *one system*, *one mp* and *two systems* benefit from PSLC pricing. Some processors running this type of configuration cannot be upgraded to meet the hardware requirements, and therefore will have higher software costs.

Configuration *two systems* has higher costs than *one system* and *one mp*, since its total capacity is likely to be higher to process the same work - see 4.1, “Effective Use of Capacity” on page 17.

In a parallel sysplex (configurations *two systems parallel* and *five systems parallel*), processor capacity can be added in smaller increments (see 5.1, “Processor Upgrade” on page 21), which minimizes software costs. Configuration *five systems parallel* can cost less than *two systems parallel* because the user can choose to run some software products on a single, small, processor.

Parallel sysplex users may be more likely to use additional software products, such as:

- CICSPlex/System Manager (CP/SM), for automatic and simplified workload management for CICS
- ESCON Manager, for simpler configuration management and fewer errors
- AOC/MVS, for operations automation, to reduce operator effort

- EPDM, for simpler consolidation of performance and capacity information across multiple processors

However, these products also provide systems management benefits in a non-parallel sysplex environment, and many non-parallel sysplex users already have them. We have not included them as differentiating factors in either the systems management or cost discussions.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	-	+	+ +

Chapter 6. Explanation of System Management Effort Items

This section explains the reasoning behind the evaluation of the system management effort lines in the table.

We have evaluated these items based on how much effort would be required for each configuration. Thus, this is the cost of the people needed to manage each configuration. Configurations requiring less effort are marked with pluses; those requiring more effort are marked with minuses.

6.1 Operational Complexity

This item address the amount of effort needed to operate the different configurations.

The more processors there are in the configuration, the more components there are to manage, and the more connections and interactions there are between them. Therefore, as the number of processors increases, the complexity and effort increase.

Configuration *one mp* is capable of being partitioned, which may require some switching of I/O devices, so it is more complex than configuration *one system*.

Parallel sysplex introduces new components, such as the coupling facility, and new function, such as data sharing, which add to the complexity. However, parallel sysplex also introduces a single point of control across multiple systems, such as:

- Console integration
- Single point of performance monitoring with RMF
- Dynamic workload balancing
- For 9672s, the Hardware Management Console, which allows multiple processor hardware to be managed as one

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	-	- -	-	-

6.2 Automated Operations

This item considers the effort in setting up and maintaining automated operations in each configuration. Automation of console operation and system activity monitoring is common for managing today′s powerful S/390 processors. IBM SystemView Automated Operations Control/MVS (AOC/MVS) is a product that provides that capability.

The single system configurations provide the simplest environment for automation, though there may be a multiplicity of subsystems, such as CICS, within the single MVS system.

The multiple system configurations are more complex. The automation must provide a focal point for automation actions, the ability to route actions to different MVS systems, a backup focal point in case the primary one fails, and more automation function to handle events unique to a multiple system configuration. However, once the automation definitions are set up, the ongoing effort of maintaining them should be no more complex than for a single system.

The parallel sysplex configurations benefit from console integration. The number of systems in the parallel sysplex should not affect the automation effort required.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	- -	-	-

6.3 Change Management

This item evaluates the effort in scheduling and implementing system changes: hardware, system software and applications. In today's competitive marketplace, the ability to respond quickly to change is an important competitive advantage.

On a single processor (configurations *one system* and *one mp*) with a single production system, changes are made once to that system. Availability needs of applications may make it difficult to find a time slot to take the system out to make changes. It may take a lot of effort to coordinate the outage between all of the people involved, including the users. If there is a problem with a change, backing it out will cause more system unavailability.

With two processors (configuration *two systems*), hardware and system software changes will need to be applied to both systems, thus requiring two change slots. They can be applied to the less critical system first, thus reducing the impact if there is a problem.

In a parallel sysplex (configurations *two systems parallel* and *five systems parallel*), processors can be removed from the parallel sysplex, changed, and then brought back in dynamically. Applications using dynamic workload balancing will be continuously available throughout this process. This allows much more flexibility in scheduling the changes, especially with five systems, where the removal of one of them still leaves 80% of the capacity available. The routing mechanism can initially route selective transactions to the upgraded system, to further test the change. Also, if there is a problem, a few transactions running on the upgraded system may fail, but then dynamic routing can route transactions away from that system. The problem can then be fixed, or the change backed out, again while keeping the application running.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	0	+ +	+ +

6.4 Problem Management

This item addresses the effort required to diagnose and fix problems.

A single system (configurations *one system* and *one mp*) provides the simplest problem diagnosis environment. However, any problem is likely to be more severe in this case, since more applications are affected. Therefore, more effort may be needed to solve it as quickly as possible.

With multiple systems (configuration *two systems*), diagnosis of some problems becomes more complex because of possible interactions between the systems. However, a problem may affect fewer applications than in the single system case.

Parallel sysplex (configurations *two systems parallel* and *five systems parallel*) introduces some additional complexity, with database data sharing, and with a transaction able to run on a different system each time it runs. However, with dynamic transaction routing in a parallel sysplex, transactions can be routed away from a system where they are failing. This results in minimal impact to the users, and thus reduces the urgency to fix the problem.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	-	0	0

6.5 Workload Management

This is the effort required to balance the workload across the available configuration, to make best use of the available processor capacity.

With only one processor (configurations *one system* and *one mp*), one MVS system, or maybe a VM/ESA system or PR/SM if test systems are also being run, controls all of the processor resources. There is no need for workload balancing across multiple processors.

With two independent MVS systems (configuration *two systems*), a system administrator must decide which workloads will run on each system, and set up the system accordingly. An example of this is shown in Figure 2 on page 37. To best use the available capacity, the administrator will need to review the setup on a regular basis, and make changes. All applications updating a common database must run on the same system, as there is no data sharing. This reduces the flexibility the administrator has in placing the systems. If a workload exceeds the capacity of a single system, the application and databases will need to be split, which could take many man-years of effort.

With a parallel sysplex (configurations *two systems parallel* and *five systems parallel*), the workloads not capable of dynamic workload balancing will need to be placed manually, as in configuration *two systems*. Even when they are not capable of dynamic workload balancing, if the databases support data sharing there is more flexibility, as these databases can be updated by applications running on different systems. The dynamic workload balancing workloads can then be set up to run in the remaining capacity. Examples of this are shown in

Figure 4 on page 39 and Figure 5 on page 39. Because these workloads are automatically scheduled by the system where capacity is available, much less effort is required in monitoring and changing the system.

If an application system grows to become larger than a single processor, then the user can either use parallel sysplex and data sharing to spread the application across multiple processors, or they must split the databases and applications. This latter option is likely to require a lot of effort, so in this scenario parallel sysplex requires much less effort in workload management.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	- -	-	-

6.6 Capacity Planning

This item is the amount of effort needed for capacity planning.

For a single processor (configurations *one mp* and *one system*), only the total capacity of all workloads needs to be planned. However, the granularity of growth may be in large increments, depending on the type of processor. The maximum size of one processor is also a limitation for some users.

For multiple processors not in a parallel sysplex (configuration *two systems*), the capacity needs of each processor must be monitored and planned separately. Capacity must be added to the processor requiring it. If the two processors must be the same size for backup reasons, then considerable effort will be required to define a workload split to optimize the use of the two processors. If the work is to be split across processors of different sizes, then they can be planned separately, but effort is required to plan how the production work can run in backup mode after a failure of one processor. The maximum size of one processor may also be a limitation in this case.

For multiple processors in a parallel sysplex (configurations *two systems parallel* and *five systems parallel*), dynamic workload balancing means that in many cases it will only be necessary to provide capacity on the basis of the total workload. Capacity can be added by upgrading any of the processors, or by adding an additional one to the parallel sysplex. This gives much greater flexibility in configuration options and in the increment of growth added. Capacity can be added as it is needed (“just-in-time”), since an outage is not required. This also allows a faster response to change, for example acquisition of another business, or sudden growth. It may also make it easier to plan for peak periods.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	- -	+ +	+ +

6.7 Performance Management

This item is the effort to monitor and tune the performance of the different systems and workloads.

For a single processor (configurations *one mp* and *one system*), there is only one system to manage, so the effort is less. However, a performance problem may impact all workloads, so effort is needed to monitor the system closely.

For multiple processors (configuration *two systems*), each system must be monitored and tuned separately. Multiple systems also introduces additional potential performance issues from contention for shared resources, for example DASD data sets and volumes.

A parallel sysplex (configurations *two systems parallel* and *five systems parallel*) simplifies the tuning effort for multiple systems. RMF's single point of reporting across the parallel sysplex reduces the monitoring effort. Also, since dynamic workload balancing takes into account the current performance of the various systems in routing transactions, any performance issues are much less likely to affect transaction response times. This reduces the urgency to identify such problems, and therefore reduces the monitoring effort required.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	0	-	+ +	+ +

6.8 Configuration Management

This item is the effort required to manage the hardware configuration, in particular the processors, the connections between them, and their connections to I/O devices.

The configuration becomes more complex as the number of processors increases. There is also a slight additional increase in complexity with the hardware required by the parallel sysplex, namely coupling facilities and sysplex timers.

A single multiprocessor (configuration *one mp*) is as complex as two separate processors (configuration *two systems*), since we assume that each side of the multiprocessor can run the workload independently if the other side should fail.

For these reasons we have rated:

One system	One mp	Two systems	Two systems parallel	Five systems parallel
0	-	-	-	- -

Appendix A. Calculation of Usable Capacity

In this section, we calculate the usable capacity of some sample configurations of each type considered in Table 1 on page 7, based on available measurements. This provides background information to the ratings given for 4.1, "Effective Use of Capacity" on page 17. We do the calculations for two different processor ranges.

Table 2 summarizes sample calculations for the IBM 9021 711-based processor range.

Table 2. Usable Capacity Calculations for IBM 9021

Configuration	**One system**	**One mp**	**Two systems**	**Two systems parallel**	**Five systems parallel**
Number of processors		1	2	2	5
Processor type		9021-962	9021-831	9021-831	9672-R61
Internal Throughput Ratio (ITR) of processor		329.08	171.79	171.79	65.31
Sum of ITRs		329.08	343.58	343.58	326.55
Data sharing impact on ITR		0	0	-14.8%	-11.3%
Workload balancing impact on ITR		0	-16.7%	0	0
Usable capacity		329.08	286.20	292.73	289.65
Effective use of capacity		95.8%	83.3%	85.2%	88.7%

Table 3 summarizes sample calculations for the IBM 9672 Model R1 processor range.

Table 3. Usable Capacity Calculations for IBM 9672

Configuration	**One system**	**One mp**	**Two systems**	**Two systems parallel**	
Number of processors	1		2	2	
Processor type	9672-R41		9672-R21	9672-R21	
Internal Throughput Ratio (ITR) of processor	51.04		29.06	29.06	
Sum of ITRs	51.04		58.12	58.12	
Data sharing impact on ITR	0		0	-10.5%	
Workload balancing impact on ITR	0		-16.7%	0	
Usable capacity	51.04		48.41	52.02	
Effective use of capacity	100%		83.3%	89.5%	

We now describe how we did the calculations.

A.1 Workload Assumptions

We assumed in 2.3, "Other Assumptions" on page 5 that 50% of the workload was using data sharing and dynamic workload balancing. In order to do these sample calculations, we need to make some more detailed assumptions about the workload running on the configurations, to give us this 50% data sharing. Our assumptions are:

- The workload is 71% CICS/IMS DB, 19% Batch and 10% TSO. We chose this as a typical workload mix for a CICS user. We chose a CICS/IMS DB workload since, at the time of writing, that was the only workload for which we had data sharing measurements.
- 70% of each CICS workload is capable of data sharing and dynamic workload balancing.
- There are four distinct workloads, which are shown in Table 4:

Table 4. Sample Workloads. Shown as a percentage of the total workload.

	Workload by subsystem			**Workload by type**		
Workload	CICS /DL1	Batch	TSO	CICS data sharing	CICS non-data sharing	Other non-data sharing
CICS 1	30			21	9	0
CICS/Batch 2	21	9		15	6	9
CICS/Batch 3	20	10		14	6	10
TSO 4			10	0	0	10
Totals	71	19	10	50	21	29

- In a non-data sharing configuration, all of a single workload must run on a single MVS system. In a data sharing configuration, all of the non-data sharing portion of a single workload must run on a single MVS system. This occurs in practice because:
  - The parts of a workload all update the same databases, and therefore must all be on the same system.
  - Separating the parts of a workload onto different MVS systems may cause an unacceptably high performance impact, due to intersystem communication or DASD contention.
  - It may be too complex operationally to split a workload across multiple systems.

A.2 Configuration/Number of Processors/Processor Type

To evaluate the usable capacity of each type of configuration, we consider specific processor models.

For the IBM 9021 processor range, we chose an IBM ES/9000 9021-962 as the single multiprocessor system. A 9021-962 can run physically partitioned, having two sides each with the capacity of a 9021-831. We chose configurations of approximately equal capacity for all the options. These are shown in Table 2 on page 33. For each configuration, the number of processors is *Number of processors*, and all of them are of type *Processor type*.

For the IBM 9672 processor range, we chose an IBM 9672-R41 as the single system, and two 9672-R21s as the two systems. See Table 3 on page 33. We chose models for which IBM had published detailed performance measurements at the time of writing - see A.3, "Internal Throughput Ratio (ITR) of Processor" on page 35.

A.3 Internal Throughput Ratio (ITR) of Processor

This is the internal throughput ratio (ITR) of one processor of the type shown in that column. Data shown here is the result of measurement runs conducted by IBM. The measurements are based on online transactions per second and batch jobs completed per hour for processors operating in an unconstrained environment. The numbers show internal throughput ratio (ITR) as measured by the IBM *Large Systems Performance Reference* (LSPR), version 4.9E, last updated in December 1994. Each processor is shown compared to the IBM 9021-520 (with an ITR of 47.5). These numbers show the capacity for processors operating under MVS/ESA SP 4.3 with the following mix of standard LSPR workloads:

- 71% - CICS
- 19% - CBW2 - a commercial batch workload
- 10% - TSO

Other control programs and workload mixes will show different capacity ratings. No allowance has been made for the impact of logical partitioning as provided by PR/SM.

The LSPR results and methodology are published in *Large Systems Performance Reference*, SC28-1187. They are also available to IBM personnel in the *LSPRPC PACKAGE on MKTTOOLS*.

A.4 Sum Of ITRs

This is the sum of the ITRs in the configuration. It is *Number of processors* multiplied by *ITR of processor*.

A.5 Data Sharing Impact on ITR

Configurations *one system* and *one mp* have no data sharing, so the impact is zero.

For configuration *two systems* there is a data sharing overhead. This comes from executing additional software instructions (for example, GRS and JES), to manage the data set level sharing. However, we have no measurements of that impact, so we cannot include that effect in our calculations.

The LSPR data that we used did not include any data sharing. To estimate the effect of data sharing on the usable capacity of configurations *two systems parallel* and *five systems parallel*, we used the methodology published in *S/390 Parallel Sysplex System Performance Considerations Presentation Guide* (MFTS PACKAGE on MKTTOOLS). Page 16 has a formula for calculating the "Parallel Sysplex Function Capacity" required to support parallel sysplex functions for both multisystems management and data sharing. We divide the whole formula by the total capacity in the configuration and derive the following formula:

```
Data Sharing Impact = M + {E + (I x (N - 2))} x P
```

where:

Data Sharing Impact = The percentage impact on the ITR from multisystems management and data sharing in a parallel sysplex
M = Multisystems management cost associated with a parallel sysplex
E = Enablement cost for data sharing in a parallel sysplex
I = Incremental growth cost for data sharing in a parallel sysplex
N = Number of MVS images participating in the parallel sysplex
P = Percentage of the total workload that is using data sharing

The multisystems management, enablement and incremental growth costs will vary based on the user's environment, the processor speed and the coupling facility processor speed.

We calculated these costs from the detailed data sharing performance measurements published in the IBM ITSO redbook, *System/390 MVS Parallel Sysplex Performance*, GG24-4356-00. These measurements used a CICS/IMS DBCTL workload (see page 52). This was a version of the LSPR CICS workload modified to use data sharing and dynamic workload balancing. The measurements used a 9674 Model C01 coupling facility. The impact of data sharing would be less with a faster coupling facility, such as a 9674 Model C02 or C03. The multisystems management cost is between 3% and 4%, so we took a value of 3.5%. We then calculated the other costs for a 9021-711 based processor from page 65, and for a 9672-R1 processor from pages 55 and 56. These values are shown in Table 5. Note that, in particular, the enablement cost may be much lower for many customer workloads, due to less locking per transaction.

Table 5. Data Sharing Cost Components

Processor type	**711-based**	**9672-R1**
Multisystems management cost (M)	3.5%	3.5%
Data sharing - Enablement cost (E)	22.5%	14.0%
Data sharing - Incremental growth cost (I)	0.5%	0.5%

We now use this formula for the different parallel sysplex configurations.

For two 9021-831s in a parallel sysplex:

```
Data Sharing Impact = 3.5% + {22.5% + (0.5% x (2 - 2))} x 50%
                    = 3.5% + 22.5% x 50%
                    = 3.5% + 11.3%
                    = 14.8%
```

For five 9672-R61s in a parallel sysplex:

```
Data Sharing Impact = 3.5% + {14.0% + (0.5% x (5 - 2))} x 50%
                    = 3.5% + 15.5% x 50%
                    = 3.5% + 7.8%
                    = 11.3%
```

For two 9672-R21s in a parallel sysplex:

```
Data Sharing Impact = 3.5% + {14.0% + (0.5% x (2 - 2))} x 50%
                    = 3.5% + 14.0% x 50%
                    = 3.5% + 7.0%
                    = 10.5%
```

These numbers are input to Table 2 on page 33 and Table 3 on page 33.

A.6 Workload Balancing Impact on ITR

The LSPR measurements all assume a single processor. Therefore, for the single processor configurations, *one system* and *one mp*, no further correction is needed. All of the workloads run on a single MVS image, all of the processor resources are accessible to all of the workloads, and the configuration can be driven up to the total capacity of the processor.

For the multiple processor configurations, we need to consider how effectively the total workload can be run across multiple systems.

A.6.1 Workload Balancing With No Parallel Sysplex

This is configuration *two systems*: two processors sharing DASD, but not in a parallel sysplex. Shared DASD means that we can place any workload on either processor. However, once a workload is placed there, it will run there until we move it manually.

In our example, we could place workloads 1 and 2 on the first system, and workloads 3 and 4 on the second system. This is shown in Figure 2. However we split up the workloads, the best case is with 60% of the total workload on one processor and 40% on the other.

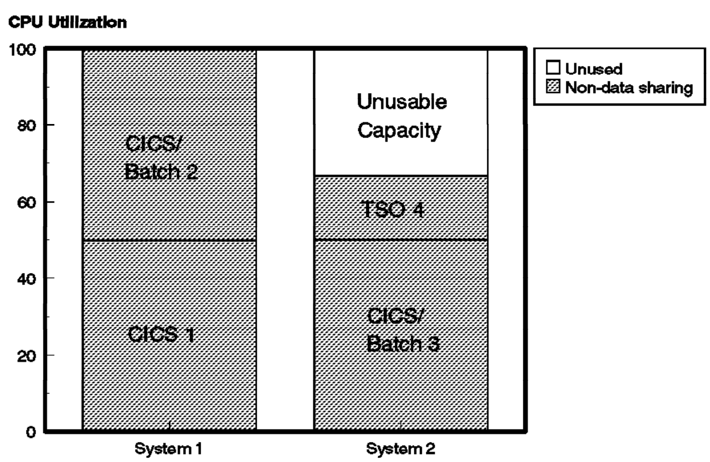

Figure 2. Workload Balancing on Two Systems With No Parallel Sysplex. When deploying the workload referenced in Table 4 on page 34, a non-parallel sysplex configuration cannot exploit the full system capacity.

Now, what happens when system 1 is full? We need to upgrade it, because we cannot split work to move it onto system 2. The overall configuration is at maximum capacity. However system 2 is only at 66.6% utilization (100% on system 1 times the workload ratio of 40/60), so the unused capacity at this time is 33.3% of a single system, or 16.7% of the total configuration. This is the value that we take into Table 2 on page 33 and Table 3 on page 33.

If you run two or more processors, you should be able to do this calculation for your own installation. Look at the last time one of the processors was upgraded, and see how much unused capacity you had on the other processors at that time.

Another way to look at this impact of workload balancing is to consider the load on each of the two processors throughout the day. Figure 3 is taken from a customer example.

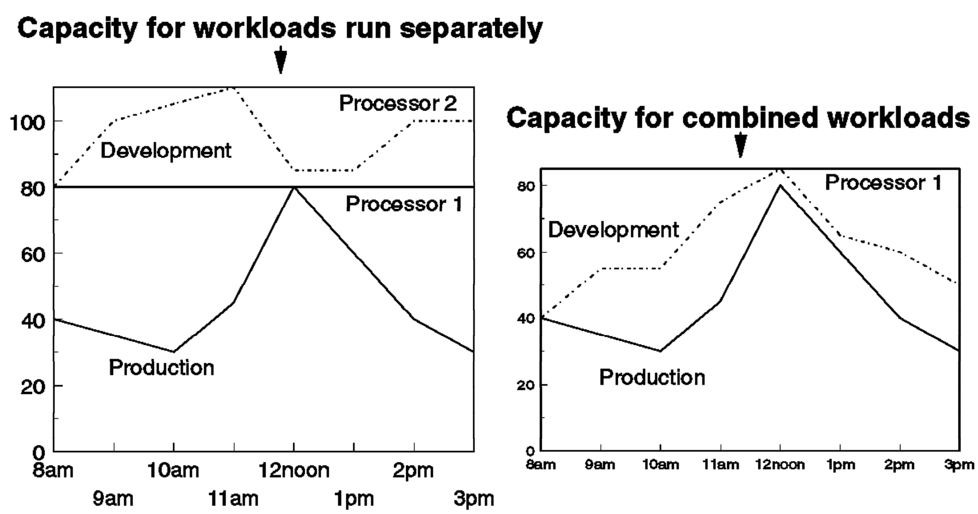

Figure 3. Two Workloads Running on Two or One Processors

With the two workloads running on separate processors, one runs the production workload, and the other runs the development workload. The two workloads have capacity peaks at different times of the day. Each processor must be large enough to handle the peak capacity of its workload, and there is a lot of unused capacity.

If both workloads are run on one processor, as in the right hand part of Figure 3, the total capacity required is less.

A.6.2 Workload Balancing in a Parallel Sysplex

In a parallel sysplex, workload can be dynamically routed across multiple MVS systems. For example, CICS transactions can be routed to any of the systems in a parallel sysplex, and can directly access databases through a database manager that supports data sharing, such as IMS DB. This was the environment measured in *System/390 MVS Parallel Sysplex Performance*, GG24-4356-00. Other transaction and database managers will support data sharing and dynamic workload balancing in the future.

Most users will not have all of their workload capable of dynamic workload balancing, and therefore we need to check whether the system can automatically route the workload to be able to use all of the available capacity in the parallel sysplex. We assumed that 50% of the workload is capable of data sharing and dynamic workload balancing. We can approach workload distribution by first mapping the non-data sharing part of each workload onto the configuration, keeping each workload on one processor. Then the data sharing part of the workload is dynamically distributed across the systems, using the rest of the available capacity.

Figure 4 shows how our sample workload could be spread across two systems in a parallel sysplex (configuration *two systems parallel*).

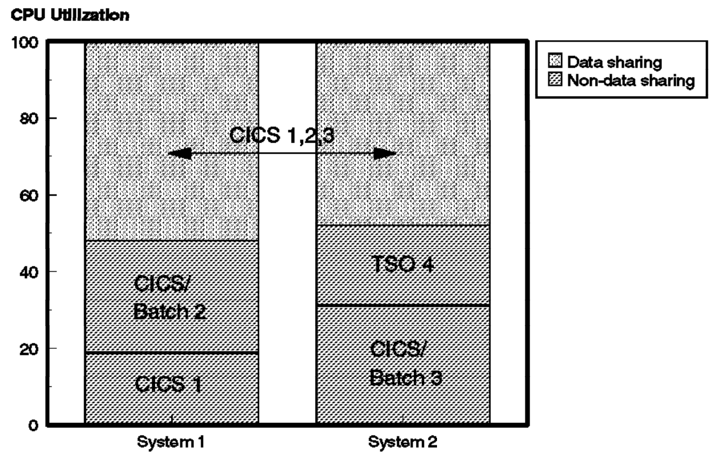

Figure 4. Workload Balancing on Two Systems in a Parallel Sysplex

Figure 5 shows how the workload could be spread across five systems in a parallel sysplex (configuration *five systems parallel*).

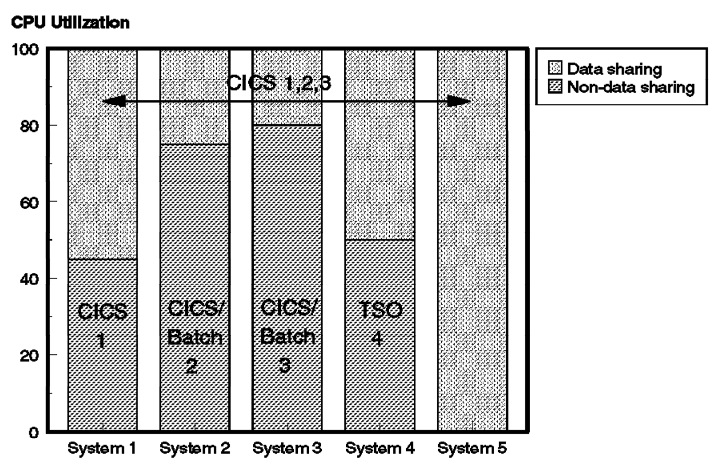

Figure 5. Workload Balancing on Five Systems in a Parallel Sysplex

As seen in the figures, in both of these parallel sysplex configurations, all of the available capacity can be used. The workload balancing impact on the total ITR is therefore zero.

A.7 Usable Capacity

This is the configuration capacity. It is the *Sum of ITRs* less the impacts of *data sharing* and *workload balancing*. This is the amount of work you could realistically process through each configuration.

A.8 Effective Use of Capacity

This shows how much of the total configuration capacity can be used. It is the *Usable Capacity* divided by the *Sum of ITRs*, expressed as a percentage. For the 9021-962, it is the usable capacity divided by the sum of the ITRs for two 9021-831s, as the 962 could be physically partitioned into the equivalent of two 831s. This therefore shows the multiprocessor effect.

A.9 Conclusions

In 4.1, "Effective Use of Capacity" on page 17, we saw that there is a cost for all forms of using multiple processor engines for a single workload. If we look at the results of our calculations in Table 2 on page 33 and Table 3 on page 33, we see that:

- A parallel sysplex may or may not be more efficient than a single system. In our example, for a 962 compared to two 831s in a parallel sysplex, the single system was better. For a 9672-R41 compared to two 9672-R21s in a parallel sysplex, the parallel sysplex was better (see the "Usable capacity" line in Table 3 on page 33). The way the two options compare depends on the multiprocessor ratios of the processor, the speed ratio between the processor and the coupling facility, and on the proportion of the workload that is data sharing.
- A parallel sysplex was more efficient than two systems sharing DASD in our example. The key factor for the shared DASD configuration is the ease of manually balancing the workloads to get the optimal split, and the effort to do that. For the parallel sysplex, the key factors are the speed ratio between the processor and the coupling facility, and the amount of data sharing workload.
- Five smaller processors was more efficient than two larger ones. The improved data sharing efficiency due to the less powerful processors more than outweighs the small cost of adding additional processors to the parallel sysplex.

Glossary

CICS. Customer Information Control System. An IBM licensed program that lets transactions entered at remote terminals to be processed concurrently by user-written application programs. It includes facilities for building, using, and maintaining data bases.

CMOS. Complementary metal-oxide semiconductor.

coupling facility. A special logical partition (LPAR) that provides high-speed caching, list processing and locking functions in a parallel sysplex.

DASD. Direct access storage device.

data sharing. In a parallel sysplex, the ability of concurrent subsystems (such as DB2 or IMS database managers) or application programs to directly access and change the same data while maintaining data integrity.

ESCON. Enterprise Systems Connection.

GRS. Global resource serialization. A function that provides an MVS serialization mechanism for resources (typically data sets) across multiple MVS images.

IBM. International Business Machines Corporation

IMS. Information Management System. A general purpose system whose full name is Information Management System/Virtual Storage (IMS/VS). It enhances the capabilities of MVS for batch processing and telecommunication and allows users to access a computer-maintained data base through remote terminals.

I/O. Input/output

I/O configuration. The collection of channel paths, control units, and I/O devices that attaches to the processor unit.

IPL. Initial program load. The initialization procedure that causes an operating system to start operation.

ITR. Internal throughput ratio. Measures the rate at which a processor can do work, relative to another processor.

ITSO. International Technical Support Organization

JES (Job Entry Subsystem). A system facility for spooling, job queuing, and managing I/O.

log off. To request that a session between a user and a computer system be terminated.

log on. To initiate a session between a user and a computer system.

mp. Multiprocessor. A processor that can be physically partitioned into two separate processors.

MVS. Multiple Virtual Storage operating system.

MVS system. A single occurrence of the MVS operating system together with its associated hardware.

parallel sysplex. A sysplex that uses one or more coupling facilities.

processor engine. A processing unit, capable of executing instructions when combined with main storage and channels.

processor. A physical collection of hardware that includes main storage, one or more *processor engines*, and channels. Also known as a central electronic complex (CEC).

PR/SM. Processor Resource/Systems Manager. A function that allows the processor to operate several system control programs (SCPs) simultaneously in logical partition (LPAR) mode. It provides for logical partitioning of the real machine.

response time. The amount of time it takes after a user presses the enter key at the terminal until the reply appears at the terminal.

RMF. Resource management facility

shared DASD. The ability for multiple MVS systems to access and update the same DASD volume. Locking is provided at the volume level by reserve/release or at the data set level by *GRS*. See *data sharing* for an extension of shared DASD.

single system image. The characteristic a product displays when multiple images of the product can be viewed and managed as one image.

S/390. System/390.

TSO. Time sharing option.

VTAM. Virtual telecommunications access method.

Index

T

U

V

W

ITSO Technical Bulletin Evaluation RED000

International Technical Support Organization
A Comparison of System/390 Configurations - Parallel and Traditional
October 1995

Publication No. SG24-4514-00

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.

ITSO Technical Bulletin Evaluation **RED000**
SG24-4514-00

Cut or Fold
Along Line

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

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM International Technical Support Organization
Mail Station P099
522 SOUTH ROAD
POUGHKEEPSIE NY
USA 12601-5400

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

Cut or Fold
Along Line

Printed in U.S.A.

SG24-4514-00