International Technical Support Organization

SG24-2502-00

System/390 MVS Parallel Sysplex Migration Paths

February 1996

IBM

International Technical Support Organization
Poughkeepsie Center

International Technical Support Organization

SG24-2502-00

System/390 MVS Parallel Sysplex Migration Paths

February 1996

Take Note!

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

First Edition (February 1996)

This edition applies to the initial version of System/390 MVS Parallel Sysplex.

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

Abstract

This document will help you select a migration path from System/370 or System/390 processors to a System/390 MVS Parallel Sysplex. It shows how you can obtain benefits from Parallel Sysplex at an early stage in the migration.

Our objective in writing this book is to enable you to identify:

- Your optimum migration path to Parallel Sysplex
- Which of your applications will benefit from exploitation of Parallel Sysplex

We consider your business requirements, and your current application workloads, hardware and software.

This document was written for technical managers and planners responsible for IBM System/370 or System/390 systems, who seek an understanding of how they can migrate to Parallel Sysplex. It is also relevant to IBM personnel who are assisting their customers to build a Parallel Sysplex migration plan. Some knowledge of System/390 and Parallel Sysplex is assumed.

(116 pages)

Contents

Abstract . . . iii

Special Notices . . . xi

Preface . . . xiii
How To Use This Book . . . xiii
Related Publications . . . xv
International Technical Support Organization Publications . . . xvi
ITSO Redbooks on the World Wide Web (WWW) . . . xvii
Acknowledgements . . . xviii

Chapter 1. Identify Your Business Requirements . . . 1
1.1 Identify Your Requirements . . . 1
1.1.1 Recording Your Requirements . . . 2
1.2 Advanced Functions for Applications . . . 2
1.2.1 Improved Application Performance . . . 3
1.2.2 Continuous Application Availability . . . 3
1.2.3 Application Growth . . . 3
1.3 System Level Benefits . . . 3
1.3.1 Improved System Performance . . . 3
1.3.2 Better Availability . . . 4
1.3.3 Capacity Growth . . . 4
1.4 Reducing the Cost of Computing . . . 4
1.4.1 Better Resource Utilization . . . 5
1.4.2 Lower Hardware Cost . . . 5
1.4.3 Lower Software Cost . . . 5
1.4.4 Higher People Productivity . . . 6

Chapter 2. How Parallel Sysplex Can Meet Your Requirements . . . 7
2.1 A Brief Introduction to Parallel Sysplex . . . 7
2.2 Aspects of Parallel Sysplex . . . 8
2.2.1 Data Sharing . . . 8
2.2.2 Dynamic Workload Balancing . . . 11
2.2.3 Single System Image . . . 12
2.2.4 Cost Effective Hardware . . . 13
2.2.5 Cost Effective Software . . . 13

Chapter 3. Understanding Parallel Sysplex Migration Options . . . 15
3.1 Basic Hardware Options . . . 15
3.1.1 Install a Coupling-Capable Processor to Replace an Existing Processor . . . 15
3.1.2 Install a Parallel Sysplex Alongside an Existing Processor . . . 16
3.1.3 Install a New CMOS Processor and Couple It to an Existing Processor . . . 16
3.1.4 Couple Existing Processors . . . 17
3.2 Application and Software Considerations . . . 18

Chapter 4. Applications . . . 19
4.1 Selecting Workloads to Exploit Parallel Sysplex . . . 20
4.1.1 Which Applications Will Run in a Parallel Sysplex? . . . 20
4.1.2 Application Selection Process . . . 20

4.2 Completing an Application Inventory . . . 21
4.2.1 Application Suite Name . . . 21
4.2.2 Business Benefit of Parallel Sysplex . . . 21
4.2.3 Risks . . . 23
4.2.4 Need ICRF or Vector Facility . . . 24
4.2.5 Total Capacity . . . 24
4.2.6 Uniprocessor Power . . . 25
4.2.7 Subsystem or Transaction Manager . . . 26
4.2.8 Database Manager . . . 26
4.2.9 Other Columns . . . 26

Chapter 5. Application Subsystem Support of Parallel Sysplex . . . 27
5.1 Exploitation of Data Sharing . . . 27
5.1.1 Data Sharing with IMS/DB . . . 28
5.1.2 Data Sharing with DB2 . . . 30
5.1.3 Data Sharing with VSAM . . . 31
5.1.4 Data Sharing with Other IBM Access Methods . . . 31
5.1.5 Data Sharing with Non-IBM Database Subsystems . . . 31
5.2 Exploitation of Dynamic Workload Balancing . . . 31
5.2.1 CICS . . . 32
5.2.2 IMS/TM . . . 38
5.2.3 Other Subsystems, Including TSO and Batch . . . 40
5.3 Priority - H/M/L/X . . . 41
5.4 Mapping Applications to Systems . . . 42
5.4.1 Isolating Workloads On a Single System . . . 42
5.4.2 Workload Distribution without Data Sharing . . . 43
5.4.3 Workload Distribution with Data Sharing . . . 44
5.4.4 Workload Distribution with Data Sharing and Dynamic Workload Balancing . . . 44

Chapter 6. Systems Management and Operational Considerations . . . 45
6.1 Systems Management and Operation of a Multisystem Environment . . . 45
6.1.1 Traditional Multisystems Considerations . . . 45
6.1.2 New Aspects of Systems Management in a Parallel Sysplex . . . 47
6.2 Backup and Recovery in a Parallel Sysplex . . . 49
6.2.1 Backup Within a Data Center . . . 49
6.2.2 Disaster Recovery . . . 51
6.2.3 Combining Backup Within a Data Center and Disaster Recovery . . . 51

Chapter 7. Software . . . 55
7.1 Software Migration Considerations . . . 55
7.1.1 Software Packaging . . . 55
7.1.2 Non-IBM Software . . . 55
7.1.3 Minimizing the Migration Time and Effort . . . 56
7.1.4 Testing the Changes . . . 56
7.1.5 Software Maintenance Strategy . . . 57
7.2 Taking Advantage of Changes in Software Pricing . . . 58
7.2.1 Graduated License Charge . . . 59
7.2.2 Parallel Sysplex Licence Charge (PSLC) . . . 59
7.2.3 Measured Usage License Charge (MULC) . . . 61
7.2.4 Indexed Monthly License Charge (IMLC) . . . 61
7.3 Base Software . . . 62
7.3.1 The Minimum Required . . . 62
7.3.2 Preparatory Work . . . 63
7.3.3 Additional Changes You Could Make . . . 66

7.4 Database Manager Software 69
7.4.1 Sharing Databases 69
7.4.2 IMS/ESA Version 5.1 Database Manager 70
7.4.3 DB2 Version 4 71
7.4.4 CICS/VSAM 71
7.4.5 Other Database Managers 71
7.5 Transaction Manager Software 71
7.5.1 CICS 72
7.5.2 IMS/ESA Transaction Manager 73
7.5.3 TSO 74
7.6 Networking and Systems Management Software 74
7.6.1 Networking Software 74
7.6.2 Systems Management Software 76

Chapter 8. Hardware 77
8.1 Core Hardware Components 77
8.2 Processor Selection 78
8.2.1 Processor Configuration Options 78
8.2.2 Processor Selection Based on Current Processors 80
8.2.3 Processor Selection Based on Application Requirements 80
8.3 Coupling Facility Selection 82
8.3.1 Coupling Facility Options 82
8.3.2 Coupling Facility Sizing 86
8.3.3 Coupling Links 86
8.3.4 Addressing Single Points of Failure 87
8.4 Sysplex Timer 88
8.5 I/O Connectivity 89
8.5.1 DASD 89
8.5.2 Tape 90
8.5.3 Network 91
8.5.4 Consoles 91

Chapter 9. Creating Your Migration Plan 93
9.1 A Sample Migration Plan 94
9.2 Next Steps 95

Appendix A. Worksheets 97
A.1 IT Requirements 97
A.2 Key Applications 98
A.3 Basic Migration Options 98
A.4 Application Inventory 99
A.5 Systems Management and Operational Considerations 100
A.5.1 Systems Management and Operation of a Multisystem Environment 100
A.5.2 Backup and Recovery in a Parallel Sysplex 100
A.6 Software Worksheets 101
A.6.1 Software Packaging 101
A.6.2 Non-IBM Software 101
A.6.3 Minimizing the Migration Time and Effort 101
A.6.4 Testing the Changes 101
A.6.5 Software Maintenance Strategy 101
A.6.6 Changes in Software Pricing 101
A.6.7 Base Software 101
A.6.8 Database Manager Software 102
A.6.9 Transaction Manager Software 102
A.6.10 Networking and Systems Management Software 102

A.7 Hardware Worksheets 103
A.7.1 Draw a Picture of Your Current Processors 103
A.7.2 Draw Pictures of Possible Hardware Changes 103
A.8 Blank Implementation Plan 104

Appendix B. Additional Material for IBM Employees 107
B.1.1 International Technical Support Organization Publications 107
B.1.2 IBM Tools Disks 107

Glossary 109

Index 111

Figures

1. Flow Through This Book . . . xiv
2. Your Business and IT Requirements . . . 1
3. 9672 To Replace an Existing Processor . . . 16
4. Noncoupled Processor with a Parallel Sysplex Alongside . . . 16
5. 9021-711 Based Processor Coupled with 9672 . . . 17
6. 9021-711 Based Processors Coupled Using 9674 Coupling Facility . . . 17
7. Typical Transaction Response Time . . . 25
8. Typical Transaction Executed on a Slower Processor . . . 26
9. IMS/TM with IMS/DB Data Sharing . . . 29
10. CICS with IMS/DB Data Sharing . . . 30
11. Examples of CICS Workload Distribution . . . 33
12. CICS with a Database Manager Not Capable of Data Sharing . . . 37
13. IMS/TM Workload Distribution without Data Sharing . . . 39
14. IMS/TM with a Database Manager Not Capable of Data Sharing . . . 40
15. Moving an Isolated Workload to a New System . . . 42
16. Combining Backup in a Data Center and Disaster Backup . . . 52
17. Exploiters of Base Sysplex (MVS/ESA Version 4) . . . 67
18. Exploiters of Parallel Sysplex (MVS/ESA Version 5) . . . 68
19. Option 1 - Install a Coupling-Capable Processor to Replace an Existing Processor . . . 78
20. Option 2 - Install a Parallel Sysplex Alongside an Existing Processor . . . 79
21. Option 3 - Install a New CMOS Processor and Couple It to an Existing Processor . . . 79
22. Option 4 - Couple Existing Processors . . . 80
23. Stand-Alone 9674 Coupling Facility . . . 83
24. Coupling Facility in LPARs on a 9672 or 9021 711-Based Processor . . . 84
25. ICMF As a Coupling Facility . . . 85
26. Backup of Coupling Facility . . . 88
27. Backup of Sysplex Timer . . . 89
28. A Logical View of Disk Connections . . . 90

Tables

1. IT Requirements . . . 2
2. IT Requirements, and How Parallel Sysplex Relates to Them . . . 8
3. Application Description and Selection for Exploiting Parallel Sysplex . . . 21
4. Worksheet: IT Requirements, and How Parallel Sysplex Relates to Them 97
5. Worksheet: Key Applications . . . 98
6. Worksheet: Application Description and Selection for Exploiting Parallel Sysplex . . . 99

Special Notices

This publication is intended to help technical managers and planners build a Parallel Sysplex migration plan. The information in this publication is not intended as the specification of any programming interfaces. See the PUBLICATIONS section of the IBM Programming Announcement for Parallel Sysplex 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:

ACF/VTAM	APPN
BatchPipes	CICS
CICS/ESA	CICS/MVS
CICSPlex	DB2
DFSMS	DFSMS/MVS
DFSMSdfp	DFSMSdss
DFSMShsm	DFSMSrmm
ES/9000	ESA/390
ESCON	IBM
IMS	IMS/ESA
MVS/DFP	MVS/ESA
MVS/SP	MVS/XA
NetView	OPC
OpenEdition	Parallel Sysplex

PR/SM	Processor Resource/Systems Manager
PSF	RACF
RAMAC	RMF
S/390	SP
Sysplex Timer	System/360
System/370	System/390
SystemView	VTAM
3090	

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

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.

Windows is a trademark of Microsoft Corporation.

Other trademarks are trademarks of their respective companies.

Preface

This document will help you select a migration path from System/370 or System/390 processors to a System/390 MVS Parallel Sysplex. It shows how you can obtain benefits from Parallel Sysplex at an early stage in the migration, and how you can achieve Parallel Sysplex data sharing within about 12 months.

Our objective in writing this book is to enable you to identify:

- Your optimum migration path to Parallel Sysplex
- Which of your applications will benefit from exploitation of Parallel Sysplex

We consider your business requirements, and your current application workloads, hardware and software.

You may discover that you have a number of options that you would like to investigate further. Your IBM representative can assist you in selecting between the different options, and in looking at the benefits and costs associated with each one.

This document is intended for technical managers and planners responsible for IBM System/370 or System/390 systems, who seek an understanding of how they can migrate to Parallel Sysplex. The document is also relevant to IBM personnel who are assisting their customers to build a Parallel Sysplex migration plan.

How To Use This Book

To build an overall migration plan, we suggest that you work through the book from the start. Each chapter provides a high level discussion first, then goes into more detail. You will be able to decide how much of each chapter you need to read as you go along.

As you go through the book you will identify some options that are of interest to you. We have provided some worksheets in Appendix A, "Worksheets" on page 97 that will help guide your thinking, and will enable you to record your thoughts for future reference. You may find it helpful to fill in these worksheets as you read each chapter.

This document is structured as shown in Figure 1 on page xiv.

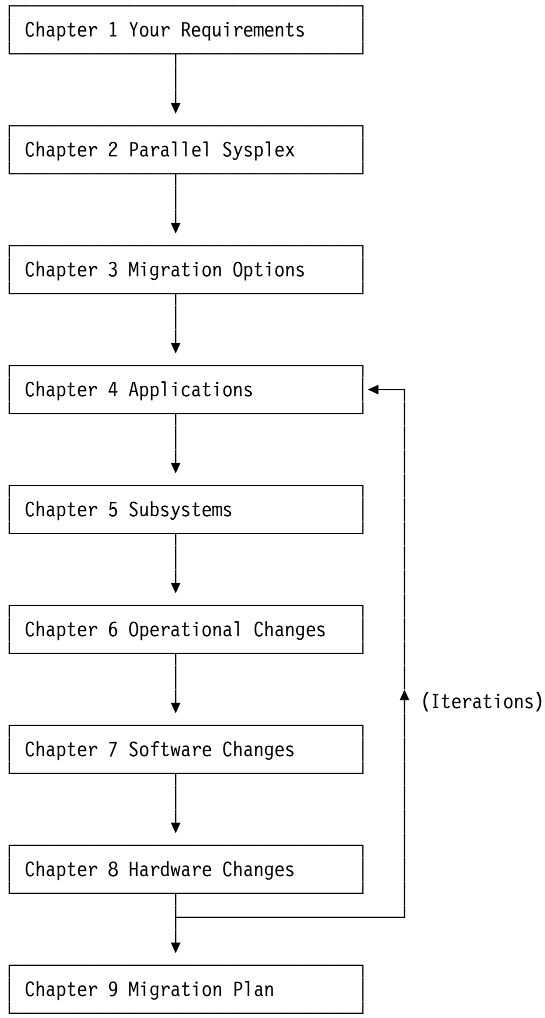

Figure 1. Flow Through This Book

- Chapter 1, "Identify Your Business Requirements"

  This will help you identify the business requirements that Parallel Sysplex can help you meet, and will help you prioritize those requirements.

- Chapter 2, "How Parallel Sysplex Can Meet Your Requirements"

  This provides a brief introduction to Parallel Sysplex, and describes the aspects of Parallel Sysplex that you may wish to implement to meet your requirements.

- Chapter 3, "Understanding Parallel Sysplex Migration Options"

  This shows you the breadth of hardware migration options available to you to implement a Parallel Sysplex. It also discusses some key software and applications aspects.

- Chapter 4, "Applications"

  This discusses how you might run your application workload in a Parallel Sysplex, and how to set up the Parallel Sysplex environment so that it suits your applications needs. It will help you decide which of your applications should exploit Parallel Sysplex first, and discusses the implications of moving some workloads to a smaller uniprocessor power.

- Chapter 5, “Application Subsystem Support of Parallel Sysplex”

  This discusses application subsystems, and their support for the Parallel Sysplex functions of data sharing and dynamic workload balancing. It also considers how you might prioritize your applications for Parallel Sysplex exploitation, and how to map them onto your processor configuration.

- Chapter 6, “Systems Management and Operational Considerations”

  This discusses operational changes you will need, or may wish, to make when you implement a Parallel Sysplex.

- Chapter 7, “Software”

  This shows you which software changes you will need for a Parallel Sysplex, and helps you decide which software can exploit the Parallel Sysplex to meet your requirements.

- Chapter 8, “Hardware”

  This shows you which hardware changes you will need for a Parallel Sysplex, and helps you select between the various options.

- Chapter 9, “Creating Your Migration Plan”

  This shows how you can combine the discussion in the previous chapters to produce your migration plan.

- Appendix A, “Worksheets”

  This contains the worksheets. You may find it helpful to complete these as you read through the chapters of this book.

- Appendix B, “Additional Material for IBM Employees”

  This supplements “Related Publications” with a list of additional publications, documents and tools available to IBM employees.

Related Publications

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

IBM employees should also refer to Appendix B, “Additional Material for IBM Employees” on page 107 for a list of additional publications, documents and tools.

- *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 JES2 Migration Notebook*, GC28-1437
- *MVS/ESA SP V5 JES3 Conversion Notebook*, GC28-1438
- *MVS/ESA SP V5 HCD: Planning*, GC28-1445
- *MVS/ESA SP V5 Conversion Notebook*, GC28-1436
- *MVS/ESA SP V5 Planning: Global Resource Serialization (GRS)*, GC28-1450
- *MVS/ESA SP V5 Planning: Install and Migration with JES2*, GC28-1428

- *MVS/ESA SP V5 Planning: Install and Migration JES3*, GC28-1429
- *MVS/ESA SP V5 Planning: Operations*, GC28-1441
- *MVS/ESA SP V5 Planning: Workload Management*, GC28-1493
- *MVS/ESA SP V5 Setting Up a Sysplex*, GC28-1449
- *BatchPipes/MVS Version 1 Release 2 Introduction*, GC28-1214
- *CICS/ESA Dynamic Transaction Routing in a CICSplex*, SC33-1012
- *IBM CICS Transaction Affinities Utility MVS/ESA User's Guide*, SC33-1159
- *VTAM Version 4.2 for MVS/ESA, VM/ESA, VSE/ESA Network Implementation Guide*, SC31-6494
- *RACF V2 R2 Planning: Installation and Migration*, GC23-3736
- *9672/9674 System Overview*, GA22-7148

International Technical Support Organization Publications

- *A Comparison of System/390 Configurations - Parallel and Traditional*, SG24-4514
- *System/390 MVS Parallel Sysplex Performance*, GG24-4356
- *MVS/ESA Sysplex Migration Guide*, GG24-3925
- *ESA/390 Microprocessor Introduction*, GG24-4497
- *DFSMS/MVS Version 1.2.0 S/390 Parallel Sysplex Support*, GG24-4400
- *RACF V2.1*, GG24-4281
- *MVS/ESA JES2 V5R1.0 and MVS/ESA JES3 V5R1.1 Technical Announcement Presentation*, GG24-3330
- *Secured Single Signon in a Client/Server Environment*, GG24-4282.

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 1-800-445-9269. Most major credit cards are accepted. Outside the USA, customers should contact their local IBM office. For guidance on ordering, send a PROFS note to BOOKSHOP at DKIBMVM1 or E-mail to bookshop@dk.ibm.com.

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

ITSO Redbooks on the World Wide Web (WWW)

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

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

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

```
http://w3.itso.ibm.com/redbooks/redbooks.html
```

Subscribing to Internet Listserver

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

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

`announce@webster.ibmlink.ibm.com`

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

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

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
```

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

Acknowledgements

This project was designed and managed by:

Brian Macfaden
International Technical Support Organization, Poughkeepsie Center

The authors of this document are:

Lynn Baker	IBM System/390 Division
Martin Ferrier	IBM UK
Choo Fatt Foo	IBM Singapore
Jim Fyffe	IBM USA
Ron Kuijper	IBM Netherlands
Gregor Neaga	IBM Germany

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

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

Ray Berard	IBM EMEA
Jose Castano	IBM System/390 Division
Ken Davies	IBM Networking Software Division
Attila Fogarasi	IBM ITSO, San Jose Center
Bob Haimowitz	IBM ITSO, Poughkeepsie Center
Alan Lee	IBM UK
David Petersen	IBM System/390 Division
Chuck Poland	IBM ITSO, Poughkeepsie Center
Nora Sweet	IBM System/390 Division

Chapter 1. Identify Your Business Requirements

Please read the "Preface" on page xiii if you have not already done so. It provides introductory material on the purpose and contents of this document, and explains how to use it.

We start this document by looking at your current requirements for information technology. It may be a requirement for more capacity, or a need to provide improved availability for a critical application. This chapter will help you identify the business requirements that Parallel Sysplex can help you meet, and will help you prioritize those requirements.

In Chapter 2, "How Parallel Sysplex Can Meet Your Requirements" on page 7 we will then show how Parallel Sysplex can meet these requirements, and therefore which aspects of Parallel Sysplex will be particularly important to you. Matching the Parallel Sysplex benefits to your requirements can also help you select the first applications to exploit Parallel Sysplex.

1.1 Identify Your Requirements

Computing systems exist to support your business. Therefore, your business will motivate your Information Technology (IT) requirements. This is shown in Figure 2.

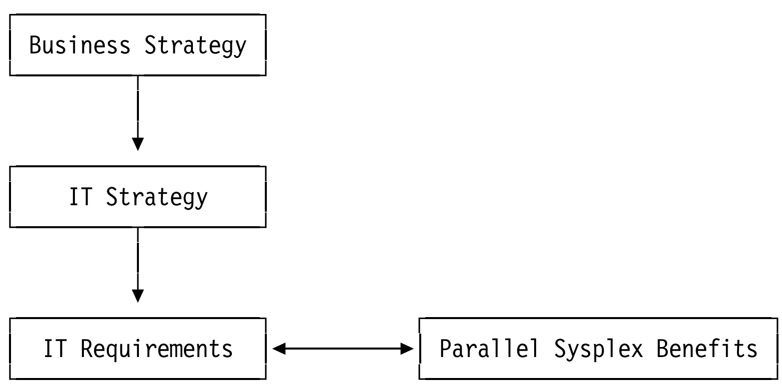

Figure 2. Your Business and IT Requirements

Your IT requirements may include:

- Implementing new business applications, or enhancing current applications, to deliver a business advantage.
- Providing the infrastructure to run current and new applications, with the required performance, availability and capacity. This also requires avoiding IT architectural constraints that would put a constraint on the business; in some cases, the maximum size of a processor or database can limit the ability to grow, or the time to do software and hardware changes can limit application availability.
- Minimizing the overall cost of the IT installation, including hardware, software, people, maintenance and environmental costs.

Parallel Sysplex can meet all of these requirements. You can achieve many of the benefits in a relatively short time, say the next 12 months.

The Parallel Sysplex benefits are discussed in Chapter 2, “How Parallel Sysplex Can Meet Your Requirements” on page 7.

1.1.1 Recording Your Requirements

We have provided worksheets in Appendix A, “Worksheets” on page 97 that relate to this chapter. We suggest that you make a copy of Table 4 and Table 5 now, and complete the worksheets as you read this chapter.

This is a list of possible IT requirements that you may have, as shown in Table 4 on page 97. As you read this chapter, complete the priority column with the relative priority for each requirement in your installation at this time. Use the instructions with the table to help you complete it.

Table 1. IT Requirements

Requirement		**Priority H/M/L**
Application	Performance	
	Availability	
	Growth	
Overall	Performance	
	Availability	
	Growth	
Cost	Resource utilization	
	Hardware	
	Software	
	People	
Additional items		

We will discuss each of these requirements in this chapter. If you identify other requirements, you can record them in the blank rows in the table.

As you read this chapter, you may identify particular key applications that have these requirements. Record these applications in Table 5.

1.2 Advanced Functions for Applications

Parallel Sysplex provides enhanced application functions that can improve application value, and thus benefit your business. We discuss here several application requirements that exist in most installations. Space is provided on Table 5 to list other requirements that are specific to your business.

1.2.1 Improved Application Performance

One of your applications may have a requirement for improved, or more consistent, performance. This may be because it is not meeting its service level agreement, or because the application exceeds the size of a single image, or because other workload interferes with it at certain times.

Parallel Sysplex allows a single application to run on multiple MVS systems, thus providing better performance. Workload can be routed to the system most able to meet the service level objective, thus providing more consistent performance and making best use of the system resources available.

1.2.2 Continuous Application Availability

Your applications may require higher availability. It is helpful to consider availability as having two components:

- High availability, avoiding unscheduled outages
- Continuous operations, avoiding scheduled outages

Parallel Sysplex helps you achieve both, by allowing an application to run across multiple systems, and thus protecting it from an outage of any one system. This includes hardware upgrades as well as changes to operating system and application software. A Parallel Sysplex can be built with various levels of redundancy, depending on your budget and application availability requirements. Availability can range from the level of availability found on today′s single processors to 7 day x 24 hour operation.

1.2.3 Application Growth

An application that is growing larger than the size of a single processor can be spread across multiple processors. Data sharing manages locking and caching between the systems. The application can continue to grow beyond a single system, and thus is not limited by the size of available hardware.

1.3 System Level Benefits

You may require:

- Improved system performance
- Better availability
- Capacity growth

across your total system. Parallel Sysplex provides an architecture to meet these requirements.

1.3.1 Improved System Performance

In 1.2.1, “Improved Application Performance” we showed how an individual application could gain improved performance from Parallel Sysplex. This can be extended to multiple applications by dynamically routing workload to the system best able to meet the required performance criteria, such that the best use is made of all the installed hardware to meet your application performance goals.

It is not a requirement that all applications use dynamic workload balancing. The ability of just a proportion of your workload to be dynamically routed may allow all applications to run with better performance. If a system is becoming overloaded, the dynamic routing of one application to a different system may allow that application to run better, and also leave more resource available on the overloaded system for the applications that cannot be dynamically routed, so that they also perform better.

Properly moving a part of your application portfolio to exploit Parallel Sysplex will improve the overall performance of your IT installation.

1.3.2 Better Availability

In 1.2.2, "Continuous Application Availability" on page 3 we showed how an individual application can gain improved availability, across both unscheduled and scheduled outages. Applications that use Parallel Sysplex data sharing can be continuously available to users. The failure of a system component, or the removal of a system component for maintenance, does not stop the application, since it is able to run on other components within the Parallel Sysplex. With the move towards network-centric computing, and the resulting rapid increase in the number of users accessing applications on your systems, high availability becomes an important requirement. The unavailability of an application can result in lost business.

1.3.3 Capacity Growth

Prior to Parallel Sysplex, your options for capacity growth were determined by the available hardware technology. You could:

- Add more capacity to a current processor, or
- Put an additional processor alongside current processors, and manually split the workload to run across more system images. This may require dividing up an application into discrete parts.

Often the technology options conflicted with the financial reality of lease termination charges, or tied you to a specific vendor's upgrade paths.

Now, with Parallel Sysplex, you still have these options, but you also have an additional alternative. An additional processor can be placed alongside current processors, and it can work as a single system image with existing systems. Therefore, Parallel Sysplex can be a more cost effective option, making better use of the new hardware investment with minimum effort. You are no longer forced to terminate leases, or take other financially unattractive options.

Growth can be tailored to your needs, with the improved granularity that the Parallel Sysplex option gives you. Additional capacity can be used for the applications that need it, but is also available to other applications as determined by the MVS workload manager.

1.4 Reducing the Cost of Computing

When we talk about the cost of computing, we do not mean just the cost of the hardware, or just the cost of the software. Rather, we mean the *total* cost of providing the computing service to your business, including hardware, software, people, maintenance, power, cooling and space.

Over many years, you have seen a continual decrease in the cost of computing for a given amount of processing. Continual hardware technology improvements, continual software developments, and continual improvements in systems management, have driven down the total computing cost. IBM is committed to continuing the reduction in the cost of computing on the S/390 platform, and Parallel Sysplex is an important part in that process. Parallel Sysplex can help you meet these requirements:

- Better resource utilization
- Lower hardware cost
- Lower software cost
- Higher people productivity

We will look at each of these requirements.

1.4.1 Better Resource Utilization

Without a Parallel Sysplex, it is likely that you would need to upgrade a processor when it is full, even though there is spare capacity on one of your other processors. The workload cannot be dynamically routed to use the spare capacity, and a static setup will never achieve perfect optimization. This is discussed in more detail in *A Comparison of System/390 Configurations - Parallel and Traditional*, SG24-4514.

Parallel Sysplex provides the capability to dynamically balance the workload across the available processors, thus making optimum use of the installed hardware. With Parallel Sysplex, you may even be able to defer an upgrade, since the spare capacity on other processors can be used. Less spare capacity to handle workload peaks is needed in a Parallel Sysplex configuration.

Note that this will also defer the associated increase in software cost, since most software is licensed according to processor capacity.

1.4.2 Lower Hardware Cost

We saw in 1.4.1, "Better Resource Utilization" that a Parallel Sysplex enables better use of current hardware, and thus delays the next upgrade, saving cost. Parallel Sysplex can also provide a more granular growth path (see 1.3.3, "Capacity Growth"), minimizing the amount of spare capacity you have and therefore again saving cost.

If you choose to implement the IBM 9672 with its CMOS technology, you will benefit from its lower acquisition and running costs (maintenance, floor space, power and cooling) compared to older technology machines. You could install a 9672 without using Parallel Sysplex, and move workload onto it in a static manner, and still gain some of the benefits of lower hardware costs. However you should carefully consider the benefits of Parallel Sysplex described in this chapter, to best meet your overall requirements.

1.4.3 Lower Software Cost

Parallel Sysplex provides a different form of software pricing, Parallel Sysplex Licence Charge, that can lower your software cost. To maximize your software savings through Parallel Sysplex Licence Charge, you should consider:

- Which of your software has Parallel Sysplex Licence Charge pricing?
- How much of your installed hardware capacity does it apply to?
- Do you need all software on all processors?

You may wish to select your hardware configuration, or change some of your software products, to maximize the savings. You may also wish to selectively run certain software (such as compilers) on a subset of your processors.

To obtain the benefits of Parallel Sysplex Licence Charge, all you need to do is to migrate your software to the correct levels, and implement a coupling facility. However, you will probably find that exploiting Parallel Sysplex for your applications will help you meet other requirements. In addition, 1.4.2, "Lower Hardware Cost" shows how exploiting a Parallel Sysplex can minimize the hardware you need, which will also minimize your software cost.

1.4.4 Higher People Productivity

Parallel Sysplex provides major new enhancements in systems management by moving towards a single system image across multiple processors. These can allow you to increase your processing capacity without increasing your people cost. A given number of people can manage an increasing amount of processing capacity, and more system users, and provide a higher quality of service at the same time. Parallel Sysplex uses your people′s existing skills and experience in IBM System/370 and System/390 hardware, software and procedures.

Chapter 2. How Parallel Sysplex Can Meet Your Requirements

This chapter describes how Parallel Sysplex can meet the requirements you identified in the previous chapter. First we provide a brief introduction to Parallel Sysplex. Then we describe the aspects of Parallel Sysplex needed to meet your requirements.

2.1 A Brief Introduction to Parallel Sysplex

Parallel Sysplex is an architectural extension to the System/390 architecture. It was announced in April 1994, 30 years after System/360 was announced, and it continues the architectural enhancements seen since System/360. Parallel Sysplex is the base for future functions, and reductions in the cost of computing, in an MVS environment. As with previous architectures, you can expect to see additional function and benefits that exploit this new architecture announced over time.

The key new feature of the Parallel Sysplex architecture is a high-speed communication method between S/390 MVS processors. This is achieved using a coupling facility. Processors are connected to the coupling facility using coupling links. The architecture is implemented in a combination of this hardware and MVS/ESA and DFSMS/MVS software. The hardware components of a Parallel Sysplex are listed in 8.1, “Core Hardware Components” on page 77 and the software components are listed in 7.2.2, “Parallel Sysplex Licence Charge (PSLC)” on page 59. Parallel Sysplex provides a very efficient method of coupling processors together. The Parallel Sysplex capacity is nearly linear up to the maximum of 32 processors. Each of the processors may have multiple tightly coupled processor engines, so Parallel Sysplex provides an efficient method of coupling many processor engines together.

In 1994, IBM also announced the 9672 Parallel Server, with higher capacity models announced in 1995. This is a new CMOS-based processor family that provides new levels of hardware price performance. IBM′s future large systems strategy is based on this industry-standard technology, and you can expect it to be enhanced as time progresses. The IBM 9672 models are coupling-capable, and therefore can participate in a Parallel Sysplex. However, they are not a prerequisite to run a Parallel Sysplex; many customers are implementing a Parallel Sysplex by linking together combinations of IBM ES/9000 9021 711-based processors and 9121 511-based processors, with or without 9672 processors.

IBM′s announcement of the Parallel Sysplex architecture is a revolution in the history of computing. A full Parallel Sysplex implementation is an extension of the S/390 architecture, and yet is very different from it, with full data sharing between processors, dynamic workload balancing, and fault tolerance. Many processors can be coupled together, all processing a common workload in a single system image. This involves many new hardware and software components and functions. As in the past, IBM′s Large Systems strategy is to deliver this new function in an evolutionary manner, avoiding the need for application change wherever possible. Most applications written for System/360 architecture, as far back as 1964, will still run on today′s hardware and software.

For more information on Parallel Sysplex see *System/390 MVS Sysplex Overview: Introducing Data Sharing and Parallelism in a Sysplex*, GC28-1208. For more information on the 9672 see *9672/9674 System Overview*, GA22-7148.

2.2 Aspects of Parallel Sysplex

In Chapter 1, "Identify Your Business Requirements," we discussed the requirements that you may have, and suggested that you complete Table 2 (a copy of Table 4 on page 97) with your priority for each requirement. At that time we considered only the first three columns of the table. The √ s in the other columns show you which aspects of Parallel Sysplex will best meet each requirement. You may wish to circle the important aspects for you on your copy of the worksheet. We will describe how these aspects deliver benefits in your requirement areas.

Table 2. IT Requirements, and How Parallel Sysplex Relates to Them

Requirement		Priority H/M/L	Data Sharing	Dynamic Workload Balancing	Single System Image	Cost Effective Hardware	Cost Effective Software
Application	Performance		√	√			
	Availability		√	√			
	Growth		√	√			
Overall	Performance		√	√			
	Availability		√	√			
	Growth		√	√	√		
Cost	Resource utilization		√	√			
	Hardware				√	√	
	Software						√
	People		√	√	√		
Additional items							

2.2.1 Data Sharing

Parallel Sysplex enables the sharing of data across multiple MVS systems. Multiple copies of an application, or multiple applications, running on different MVS systems, can access and update a common database.

First we explain what data sharing is. Then we show how it provides benefits related to your requirements.

2.2.1.1 Explanation of Data Sharing

Whenever data is shared between two or more systems, an access protocol is required to ensure that the systems do not update the same piece of data at the same time, thereby making it invalid. Before writing or updating data, a system must obtain a lock associated with that data. While one system holds the lock, no other system can obtain it.

For many years, IBM S/390 systems have been able to share data at a *data set level*, using shared DASD. Each MVS system can access and update the data on the DASD volume. Before updating any data, a system obtains a lock, either at the DASD volume level using reserve/release, or at data set level using Global Resource Serialization (GRS). While the update is in progress, no other system can update any data on that volume (for reserve/release) or data set (for GRS). This implementation is adequate for data sets that are infrequently updated, but is clearly not satisfactory when multiple systems need to do a lot of updates to records in a database.

Parallel Sysplex extends shared DASD by using a coupling facility as a high-speed store for locks and data buffers. These facilities are used to provide *record level data sharing*. Different MVS systems can update the same database at the same time; locks are held by record, or groups of records.

The coupling facility also provides a highly efficient means for maintaining cache coherency across the data sharing systems. This mechanism enables each system to continue to use local processor memory caching for maximum data access efficiency. Coherency of data in processor memory is maintained using a hardware assisted cross-invalidate mechanism. When an update to shared data is made from one node in the sysplex, the coupling facility sends a cross-invalidate signal to any other systems that are caching the same granular data (record) in memory. This signal is performed by the coupling link hardware, without any processor interruption and without any operating system software intervention. It is this mechanism that, coupled with the hardware assisted global locking mechanism, provides for the nearly linear scalability offered by the S/390 Parallel Sysplex.

2.2.1.2 Benefits of Data Sharing

Data sharing allows applications on different MVS systems to update the same data. This clearly provides more flexibility in placing applications across multiple systems. It also allows a single application to be split into multiple parts, with those parts running on different MVS systems. We will describe the benefits that this provides, related to the requirements listed earlier. Note that additional benefits come from using dynamic workload balancing with data sharing; see 2.2.2, “Dynamic Workload Balancing” on page 11 for details of those benefits.

Performance: An individual application may be performance constrained because it is exceeding the resources available on one system. Parallel Sysplex enables an application to be spread across multiple systems, thus providing more resources for it to use. This greater parallelism can result in improved performance.

At a system level, applications may be contending for system resources, resulting in degraded performance for some applications. If some applications can be split across multiple systems, using data sharing, all applications can benefit from a reduction in contention.

Availability: If multiple copies of an application are run on multiple processors then the failure of a subsystem, system or processor will only affect a subset of the copies. Some copies will remain available, and users can immediately log on again to the available systems. (Note that this is significantly enhanced if dynamic workload balancing is also used - see 2.2.2, "Dynamic Workload Balancing" on page 11).

Even if only one copy of an application is normally run, the use of data sharing can reduce application outage time. The failure of a system will cause the application to fail. However, if the databases for that application are using data sharing, recovery of the application on an alternative system can be much faster, since lock information can be held in a coupling facility external to any processor. Thus, applications can be started on an alternative system without the need to resolve all of the held locks first.

Perhaps more importantly, application availability can be improved when a subsystem, system or processor needs to be changed or upgraded (scheduled outages). Without data sharing, an application runs on only one system (if it updates databases). To take down that system for maintenance, all users must log off, then the application must be closed down, and then it can be restarted on an alternative system. With data sharing, the application can run on both systems at the same time. While users may still need to log off the system being taken down for maintenance, and log on to the alternative system, each user can now do this at their convenience. The application remains continuously available to the users.

Thus, data sharing allows you to improve the availability of your applications.

Growth: When an individual application exceeds the size of a single processor, then either the processor must be upgraded, or the application must be split to run across multiple systems. Data sharing makes the splitting of an application much easier.

Without data sharing, only one of the processors could update each database, so careful splitting of the application would be required. The application and database would need to be split in a consistent manner, and this could take many people-years of effort. Data sharing allows the various pieces of the application, on several systems, to update one set of databases with integrity. This type of split is much easier to achieve, and the application can be allowed to grow.

At a system level, once a processor is full then either the processor must be upgraded, or the workload must be split to run across multiple systems. Again, data sharing makes the splitting of the workload more feasible, since applications on different processors can update a common database. This provides additional flexibility in selecting additional capacity, since additional processors can be added in parallel to current ones without a need for a careful workload and data split. Also, since Parallel Sysplex allows an application to be spread across multiple processors, the individual processors no longer need to be larger than your largest application.

Resource Utilization: Because applications and workloads can be split to run across processors, there is much greater flexibility in which workloads run on which systems. This allows better use of the hardware resources that you have. Also, because adding new systems is so much simpler, hardware can be purchased in smaller increments, as required, thus reducing the unused capacity that you have.

People: We have already seen that data sharing reduces the complexity of many tasks, including performance management, availability management, and workload management to allow growth. This reduces the amount of work in systems management, and makes your people more productive.

2.2.2 Dynamic Workload Balancing

Dynamic workload balancing builds on data sharing. Dynamic workload balancing is currently implemented in CICS/ESA, and is planned to be implemented in IMS/ESA also. With dynamic workload balancing, an incoming transaction is dynamically routed to the system that is best able to complete that work at that time. Thus, if multiple copies of an application are run on multiple processors, the system can automatically route work between those processors, based on the service level objective for that work. Dynamic workload balancing requires data sharing, since multiple processors must be able to update a common database. We describe the benefits that this provides, related to the requirements listed earlier.

Performance: With dynamic workload balancing, the performance of the systems is monitored, and transactions are routed to the system best able to meet the service level objective at that time. Thus, for an application running on multiple systems, a performance problem on one system will have little impact, since transactions will be dynamically routed to other systems.

At a system level, if some of the applications use dynamic workload balancing, all applications can benefit. If one system is overloaded, applications using dynamic workload balancing will route their work to other systems, thus leaving resources available for workloads that can only run on that system. For example, a large IDMS data base (not sharing data) must run on a single processor. This processor can also run a CICS/DB2 workload that uses dynamic workload balancing. During peak IDMS work, CICS will route the DB2 workload to other processors in the Parallel Sysplex. When IDMS work is less, the DB2 workload can make use of this processor.

Availability: With dynamic workload balancing, the availability of the systems is also monitored. For an application running across multiple systems, if any needed component of a system is not available, work will be dynamically routed to other systems instead. Thus the loss of a system does not mean the loss of the application to the end users. When a failure occurs, only the transactions actually in progress at the time of the failure need to be recovered. All incoming transactions can be automatically routed to the systems that are still operating, so the application is still available. Once the failing system has been restarted, work will be routed to it again. If the user happens to be connected to a system that fails, then they will need to log on again. They can do that immediately, and with VTAM generic resources they will be automatically logged on to a copy of the application that is still running.

Perhaps more importantly, the application can remain available when a subsystem, system or processor needs to be changed or upgraded (scheduled outage). As the system is removed, incoming transactions can be automatically routed to the remaining systems, and processed as normal. When the system is reintroduced, work can be routed to it again. The user will see no outage, unless they are connected to the system being upgraded. In that case they will need to log off, and log on again to a different system. They can do that at a time convenient to them.

Thus, dynamic workload balancing with data sharing allows you to improve the availability of your applications

Growth: With dynamic workload balancing, transactions from an application can be dynamically routed across multiple systems. This allows an application to grow across multiple systems, while still appearing as a single application to the users.

Resource Utilization: As we have seen, dynamic workload balancing routes work to where the resources to process it are available. Thus, it provides better use of available hardware resources. This can minimize the amount of spare hardware capacity required.

People: Dynamic workload balancing further reduces the complexity of many tasks, including performance management, availability management, and workload management to allow growth, since the system is able to dynamically route work to achieve the service level objective. This reduces the amount of work in systems management, and makes your people more productive.

2.2.3 Single System Image

Parallel Sysplex allows the management of multiple MVS systems as a single system. It also allows the use of a single point of control. This simplifies many systems management tasks, including:

- Workload management - dynamic workload balancing avoids the need for manual effort.
- Change management - Parallel Sysplex provides the ability to remove and return a system to the Parallel Sysplex nondisruptively. This allows system maintenance to be scheduled at more acceptable times, and thus reduce overtime costs. It will certainly reduce effort in change scheduling.
- Operations management - a Parallel Sysplex can be operated from a single operations console, and commands can be automatically propagated to multiple systems, including vary offline of devices and RACF commands.
- Console automation is simplified with the single point of control.
- Performance monitoring and management is simpler, with a single point for performance reporting. Accounting is also simplified.
- Availability management - a Parallel Sysplex has built in redundancy, since applications can continue to run despite the loss of a processor, MVS system, or subsystem. Recovery is also less urgent, and simpler.
- Cloning of systems and subsystems - with shared MVS software definitions and libraries on DASD. Multiple MVS systems can IPL from the same SYSRES, thus minimizing the amount of DASD needed and reducing the effort in managing multiple copies.
- Databases are accessible to all systems, eliminating the need to partition database access.

Growth: The single system image of a Parallel Sysplex means that extra processors can be added much more easily. This improves flexibility, allows you to take advantage of new technology, and enables you to install capacity as you need it, rather than when a major outage can be scheduled.

Hardware Cost: Because the single system image gives you improved flexibility in selecting hardware and enables you to add and remove capacity as you need it, you can optimize your hardware costs without increasing your people costs.

People: With a single system image, a given number of people can manage an increasing amount of processing capacity, and more system users, and provide a higher quality of service at the same time. Parallel Sysplex uses your people′s existing skills and experience in IBM System/370 and System/390 hardware, software and procedures. The result is lower people cost for a given amount of processing capacity.

2.2.4 Cost Effective Hardware

The IBM 9672 CMOS-based processors cost much less to run than older technology processors, in the following areas:

- Lower hardware maintenance cost
- Lower power and cooling requirements
- Lower floor space requirements
- Improved granularity in capacity growth, as they are available in smaller growth increments

If you want to reduce your hardware costs, the IBM 9672 can help.

Parallel Sysplex allows you to implement multiple processors, either multiple 9672s or 9672s alongside bipolar technology processors, more effectively, by retaining a single system image and by balancing your workload across multiple systems. Also, because adding new systems is so much simpler, hardware can be purchased in smaller increments, as required, providing greater flexibility and lower cost.

2.2.5 Cost Effective Software

Parallel Sysplex introduces a different form of software pricing called Parallel Sysplex Licence Charge. This is an attractive pricing option for customers with a Parallel Sysplex. It can reduce the software cost dramatically.

In addition, you may have some software that only needs to run on a subset of the processors: for example, compilers, the JES/328X program for printing, or some systems management products. If these are moved onto an IBM 9672 they can benefit from a lower capacity for software charging.

These are discussed more fully in 7.2, “Taking Advantage of Changes in Software Pricing” on page 58.

Chapter 3. Understanding Parallel Sysplex Migration Options

This chapter shows you the breadth of hardware migration options available to you to implement a Parallel Sysplex. We also discuss some key software and applications aspects of Parallel Sysplex. These areas are discussed in more detail later, but it is helpful to understand the breadth of your options before we discuss your applications.

3.1 Basic Hardware Options

To implement a Parallel Sysplex, you must have a coupling facility, and at least two processors that are capable of connecting to it. To provide a coupling facility, you can use a dedicated machine (such as an IBM 9674 Coupling Facility), or a PR/SM LPAR in certain processors. There is also a unique set of Parallel Sysplex Licence Charge qualification criteria for a single system Parallel Sysplex using Integrated Coupling Migration Facility (ICMF). We will not discuss the coupling facility options here; they are covered in 8.3, “Coupling Facility Selection” on page 82. For now, just note that there are several ways to implement a Parallel Sysplex.

We will call a processor that is able to connect to a coupling facility a *coupling-capable processor*. At the time of writing, the coupling-capable processors available from IBM are:

- IBM 9672 processors
- IBM ES/9000 9121 511-based processors
- IBM ES/9000 9021 711-based processors

We discuss the processor options in 8.2, “Processor Selection” on page 78. Just note that both current range ES/9000 processors and the new CMOS based 9672 are coupling capable, giving you a range of options.

Whatever your current hardware configuration, you have several basic hardware migration routes to Parallel Sysplex. Some, or all, of these will be feasible in your installation. We list the options, and include just one example of each option. As we go through the options, you may like to record the ones applicable to you on the worksheet A.3, “Basic Migration Options” on page 98.

3.1.1 Install a Coupling-Capable Processor to Replace an Existing Processor

If you have an existing processor that is not coupling capable, you could choose to replace it with a coupling-capable processor. An example is shown in Figure 3 on page 16. This could be a 9672 Model R with a single CEC and ICMF. It may only run one MVS system in the Parallel Sysplex; that MVS system will be connected to ICMF as a *monoplex*. See 3.2, “Application and Software Considerations” on page 18 for benefits of this environment.

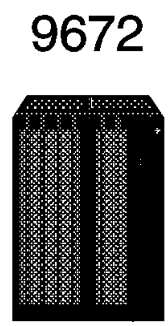

Figure 3. 9672 To Replace an Existing Processor

3.1.2 Install a Parallel Sysplex Alongside an Existing Processor

If you have an existing processor that is not coupling capable, or that does not currently have the software to support Parallel Sysplex, you could choose to install a Parallel Sysplex alongside your existing processor, but not coupled to it. An example is shown in Figure 4. This consists of two or more 9672s, coupled using a coupling facility. Note that the coupling facility could be in an LPAR on one of the 9672s. For future capacity upgrades you can add additional coupling-capable processors, and eventually replace the existing processor. Note also that, since certain single coupling-capable processors running ICMF qualify as a Parallel Sysplex, this option could be implemented by adding just one processor.

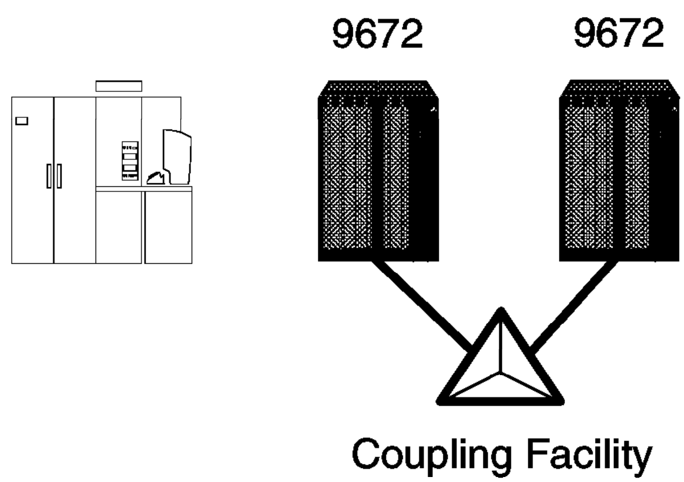

Figure 4. Noncoupled Processor with a Parallel Sysplex Alongside

3.1.3 Install a New CMOS Processor and Couple It to an Existing Processor

If you have an existing processor that is coupling capable, or that can be upgraded to be coupling capable, then you could choose to install a new coupling-capable processor and couple it to your existing processor. An example is shown in Figure 5 on page 17. This is a 9672 coupled to a 9021-711 based machine, using an IBM 9674 Coupling Facility.

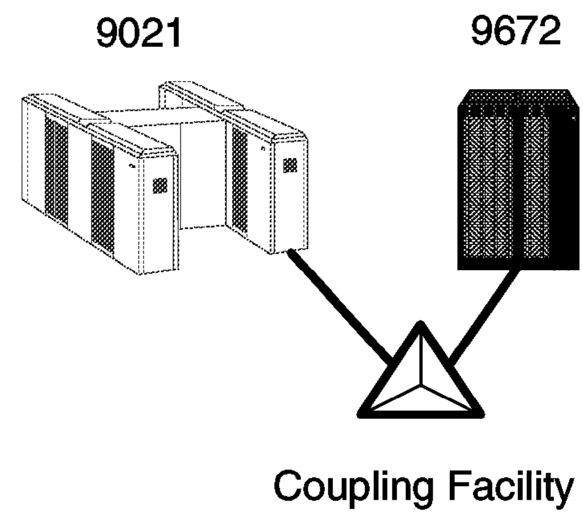

Figure 5. 9021-711 Based Processor Coupled with 9672

3.1.4 Couple Existing Processors

If you have multiple existing processors that are coupling capable, or that can be upgraded to be coupling capable, then you could couple them together in a Parallel Sysplex. An example is shown in Figure 6.

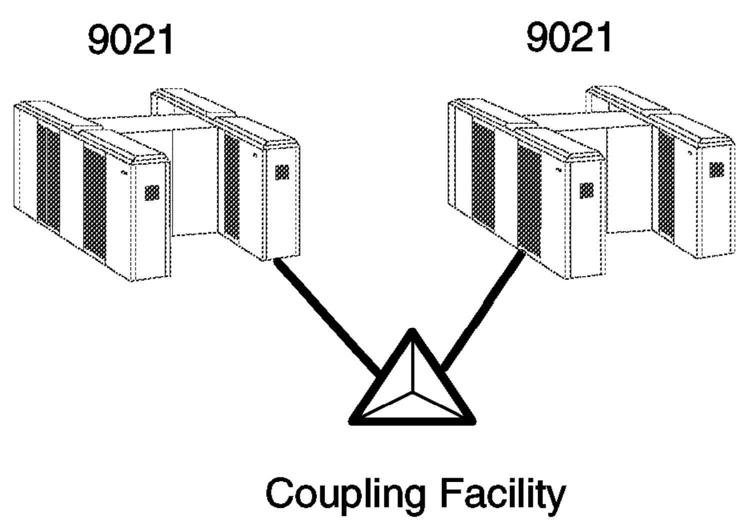

Figure 6. 9021-711 Based Processors Coupled Using 9674 Coupling Facility

3.2 Application and Software Considerations

While considering Parallel Sysplex migration, you should keep in mind the following points:

All System/390 Applications Will Run in a Parallel Sysplex Environment: That is to say, they can be moved onto MVS/ESA Version 5.1 systems that are coupled, and they will continue to run as before. The only thing you need to check is that the software the application uses is supported under MVS/ESA Version 5.1.

Without a New Processor, the Only Change is the Parallel Sysplex Itself: If you are not introducing a new processor type, then the only change you need to consider is the Parallel Sysplex itself. If you are introducing a new processor type, such as an IBM 9672, then you need to consider differences in operational characteristics. Also, if the new processor type has less capacity than your current machines (total capacity or uniprocessor power), then you will need to consider the effect of that on your workload. Topics 4.2.5, "Total Capacity" on page 24 and 4.2.6, "Uniprocessor Power" on page 25 show you how to do that.

Parallel Sysplex Is Relevant to a Single MVS System: Parallel Sysplex is also relevant if one of your key applications must run on a single system because it does not yet exploit data sharing. You can retain the single system on any of the coupling-capable processors, and implement a single system Parallel Sysplex. This will give you the benefits of the MVS System Logger and Parallel Sysplex Licence Charge Extended software pricing. Then, when you require more capacity, or when your application is able to data share, you will have two options for further capacity growth: upgrade the single system to a larger one, or add an additional processor to the Parallel Sysplex, which will preserve the single system image of Parallel Sysplex. This second option will allow you to gain more of the benefits discussed in Chapter 2, "How Parallel Sysplex Can Meet Your Requirements" on page 7.

Most Workloads Can Run on the IBM 9672: Some of the sample hardware scenarios include the installation of an IBM 9672. While the 9672 was originally announced in many countries as a *transaction* server, most System/390 workloads will run on it, including:

- Online transaction processing, such as IMS and CICS
- TSO
- Batch
- Testing

We discuss the movement of workload onto a 9672 in more detail in Chapter 4, "Applications" on page 19.

Chapter 4. Applications

In Chapter 1, “Identify Your Business Requirements” and Chapter 2, “How Parallel Sysplex Can Meet Your Requirements” you should have identified your IT requirements, and seen how Parallel Sysplex can help you meet them. This chapter illustrates how to select the first workloads that you wish to exploit Parallel Sysplex.

If your requirements are system wide, this chapter will help you to prioritize which workloads you wish to exploit Parallel Sysplex. For example, using dynamic workload balancing for one application may allow much better use of your current hardware resources. If you identified one or more key applications in Chapter 1, “Identify Your Business Requirements,” and perhaps recorded them in Table 5 on page 98, then this chapter will help you decide how those applications can exploit Parallel Sysplex.

If you initially intend to only exploit system related Parallel Sysplex functions, such as RACF, GRS, JES, or NetView, then you may skip this chapter and proceed with Chapter 6, “Systems Management and Operational Considerations” on page 45.

Many large systems environments will not migrate all their workload to a Parallel Sysplex at once. Initially, your Parallel Sysplex might include only a selected subset of your processors and applications. The key questions you need to answer are:

- Which of your applications will run in a Parallel Sysplex?
- Which ones can exploit Parallel Sysplex functions?
- Which ones you will migrate first?

This chapter looks at the selection of workloads to exploit Parallel Sysplex. We will look at the business requirements and risks that affect the sequence in which you migrate applications to a Parallel Sysplex environment. We will also discuss workload capacity requirements.

In the next chapter we will look at application subsystems, and how they enable applications to use data sharing and dynamic workload balancing. At the time of writing, application subsystems (for example, CICS, IMS, and TSO) support Parallel Sysplex related functions, such as data sharing and dynamic workload balancing, to different degrees. Although these systems will evolve over time, knowing the currently available functions of each subsystem will help you put together your individual migration roadmap.

As you go through this chapter, you will learn about the key application related aspects of Parallel Sysplex migration. You will eventually need to investigate your applications and collect the data required for application selection. Table 3 on page 21 is an example that can help you in this process (a blank copy is provided in Table 6 on page 99 for use as a worksheet). By the end of the next chapter, you should be able to prioritize your applications, using this table.

4.1 Selecting Workloads to Exploit Parallel Sysplex

You may or may not decide to turn your computing environment into a Parallel Sysplex all at once. As with any major architecture change in the past, starting with a partial implementation can make sense. Think back to when you introduced DFSMS, for instance. You had only a subset of your disks SMS managed, and you expanded that subset over time. Or consider the way you installed ESCON. You might have used fiber optic connections for only some of your peripherals in the beginning. In each case, substantial benefits could be derived from a partial implementation of the new architecture in the early stage.

Whether you convert some or all of your computing environment to a Parallel Sysplex, there is a need to evaluate which subset of applications or workloads you want to migrate to exploit the Parallel Sysplex.

4.1.1 Which Applications Will Run in a Parallel Sysplex?

This point may be obvious to some, but is included here for the sake of completeness. It is important to recognize that *every* S/390 application is Parallel Sysplex compatible in the sense that it can run on a Parallel Sysplex *without any changes*, provided only that the software supporting the application runs on MVS Version 5.1 or later. You can turn your computing environment into a Parallel Sysplex and run all your applications on it from day one, as-is.

Now, whether an application can exploit specific Parallel Sysplex functions such as data sharing and dynamic workload balancing is another question that we will examine through the rest of this chapter and the next.

4.1.2 Application Selection Process

To decide which applications you want to move to Parallel Sysplex first, we suggest you go through a formal application selection process. A good way of doing this would be as follows:

1. Establish firm requirements.

   Based on Chapter 1, "Identify Your Business Requirements," ask yourself what it is that your IT environment and your applications need most. Is it, for example, reduced software and hardware cost, is it capacity growth, or is it continuous availability? Only a firmly defined requirement list will enable you to decide on your optimal application migration sequence.

2. Build your list of ranking criteria.

   You need to come up with a list of weighted criteria to assess applications with respect to their suitability for Parallel Sysplex exploitation. These criteria will include aspects like:

   - The applications′ business relevance
   - Any business requirements that demand the application′s exploitation of Parallel Sysplex
   - The risks involved with the migration
   - The capacity, performance, throughput and elapsed time requirements
   - The Parallel Sysplex related functions that the underlying application subsystem supports

3. Build an inventory of applications or application suites.

   You will need to collect, for each application, various information relating to application migration planning, including data base systems and transaction monitors used as well as capacity and performance requirements. It may not always be feasible to look at every single application. Rather, you should define application suites. That is, groups of applications with, for instance, related functions or common data or a common subsystem platform. A sample application inventory is shown in Table 3.

4. Rank each application suite using your criteria.

   By applying your ranking criteria to your application inventory, you should now be able to select the applications to migrate first. Furthermore, you should be able to define your overall long-term optimum migration sequence.

4.2 Completing an Application Inventory

Table 3 contains a sample application inventory table. A blank copy of this table can be found in Table 6 on page 99, and you may wish to fill it in for your applications. We will explain the columns in this table as we proceed through the sections of this chapter and the next one. Each column contains a reference to the relevant sections. On your first pass through you may not have all of the information asked for readily available. You may still gain value if you use estimates where appropriate, as the table will then point to your key applications, which you can then investigate further.

Table 3. Application Description and Selection for Exploiting Parallel Sysplex

Application Suite Name (4.2.1)	**Business Benefit of Parallel Sysplex (4.2.2)**	**Risks (4.2.3)**	**Need ICRF or Vector Facility (4.2.4)**	**Total Capacity (4.2.5)**	**Uniprocessor Power (4.2.6)**	**Subsystem or Transaction Manager (4.2.7)**	**Database Manager (4.2.8)**	**Exploitation of Data Sharing (5.1)**	**Exploitation of Dynamic Workload Balancing (5.2)**	**Priority - H/M/L/X (5.3)**
Order Entry and Invoice	H	M	Y	80	20	CICS	VSAM, IMS/DB	F Y	Y	L
Warehouse Stock	M	L	N	50	5	CICS	DB2	Y	Y	M
Payroll	L	H	N	80	10	Batch	VSAM	F	N	X

4.2.1 Application Suite Name

In this column, enter your application names. You will probably want to list your major applications, and then group the other applications by related functions, use of common data, and common subsystems.

4.2.2 Business Benefit of Parallel Sysplex

Even if your Parallel Sysplex migration is primarily motivated by potential software and hardware cost reductions, you want to look at the business relevance of the applications being migrated first. In this section, we discuss these application related aspects.

By the end of this section, you should be able to complete this column of Table 6 on page 99 with the business benefit that this application could gain from exploiting Parallel Sysplex. We suggest that you use:

H if this application would provide a large business benefit
M for a medium business benefit
L for a low business benefit

We discussed the business benefits of Parallel Sysplex to your applications in 1.2, "Advanced Functions for Applications" on page 2. Some further suggestions on classifying your applications are given below:

- Providing higher application availability

  High availability is the ability to avoid or mask unplanned outages of a computer system. As discussed in 2.2.2, "Dynamic Workload Balancing" on page 11, Parallel Sysplex allows exploiting applications to be virtually unaffected by any unplanned outages.

  High availability can be of utmost importance for business applications processing high money values, like banking applications. It is equally important where high transaction rates occur with a direct business relationship, such as point-of-sale systems. Such applications might be among the first to be considered to exploit Parallel Sysplex.

- Providing continuous operation

  Continuous operation is the ability to run a computer system (in most cases, an online transaction service) day and night, all year, without any planned or unplanned outages. Planned outages of online transaction services can be caused by batch jobs requiring exclusive use of the data bases. Usually, this can only be solved through appropriate application redesign. However many planned outages are a result of the need to regularly schedule a system outage for software or hardware maintenance. We saw in 2.2.2, "Dynamic Workload Balancing" on page 11 that if multiple systems are available in a Parallel Sysplex, planned outages of a system can be masked by dynamically redirecting workload to other systems.

  Continuous operation is required where a service needs to run day and night, such as communications networks, hospital services, and again, banking. It is also required for applications that are available world-wide, or applications that service a network of users spanning multiple time zones.

- Providing capacity growth

  Some business applications, especially in the period after their introduction, tend to rapidly and repeatedly outgrow the capacity of a given system. In a conventional configuration, it is not technically feasible to add processor resources in small increments. As a result, many IT organizations decide to initially install large processor capacities for growing workloads. These capacities might be under-utilized in the beginning, but will fill up during their depreciation period.

  This approach, though typical, does not achieve the best value for money. A processor that is only used, say, 25% at the beginning, and 75% at the end of a given time period, has only had an average usage of 50% during this time. Simply stated, you could save 50% of the cost if you could install these processing capacities in multiple, small increments.

  As discussed in 2.2.2, "Dynamic Workload Balancing" on page 11, a Parallel Sysplex configuration can take advantage of small processor capacity increments, because it has the potential to spread workload over multiple processors. If you have an application that may grow substantially in the future, it might be a very good candidate to exploit your Parallel Sysplex environment.

- Migrating applications with long term relevance

  Applications with a long term relevance to the business will give greater total benefits from exploiting Parallel Sysplex. Many applications can exploit Parallel Sysplex without change, but if modifications to an application are needed, make sure the effort is applied to an application with a long expected life.

  Any application can run in a Parallel Sysplex with no migration effort, as long as it does not require any Parallel Sysplex specific functions such as data sharing and dynamic workload balancing. Applications with a limited lifetime should be migrated to Parallel Sysplex without change.

Having read this section, you can now apply it to your applications and fill out the column labelled *Business Benefit of Parallel Sysplex* in Table 6 on page 99.

4.2.3 Risks

By the end of this section, you should be able to complete this column of Table 6 on page 99 with the business risk involved in this application exploiting Parallel Sysplex. We suggest that you use:

H if the business risk is high
M if the business risk is medium
L if the business risk is low

Besides the benefits that an application can get from moving to Parallel Sysplex, there are also potential risks that need to be considered. Typically, a Parallel Sysplex configuration contains new hardware components. Initially there are unfamiliar operating procedures, and the application subsystems may use new functions with which the database administrators are not yet familiar. All of these add up to risks for application migration to Parallel Sysplex, especially in the early stages. These risks should be understood and managed.

If the first application to exploit Parallel Sysplex could have a major business impact, the migration needs to be carefully planned. Some suggestions to minimize risk are:

- Migrate less critical applications first

  To become experienced with the Parallel Sysplex migration, it makes sense to tackle some less critical applications first. In the course of doing that, the IT organization can become familiar with new hardware elements, operating procedures, and subsystem software components.

  If you decide to migrate an application with high business impact relatively early, it is essential that it goes through a solid testing period. An isolated Parallel Sysplex on multiple 9672s, or using PR/SM LPARs, is an ideal testing environment for that purpose.

- Exploit Parallel Sysplex for new applications

  An attractive alternative is to develop new applications so that they exploit Parallel Sysplex from the start. Development is inherently less prone to risk than migration, since extensive testing of the new application is already planned. When developing applications with Parallel Sysplex in mind, you can choose the subsystems best suited for Parallel Sysplex support. You can easily avoid workload affinities and plan for data sharing and dynamic workload balancing. New applications are also likely to last a long time, so the long term business value of Parallel Sysplex exploitation can be high.

Having read this section, you can now apply it your applications and fill out the column labelled *Risks* in Table 6 on page 99.

4.2.4 Need ICRF or Vector Facility

Vector Facilities and Integrated Cryptographic Features (ICRF) are not available on current models of the IBM 9672. A Parallel Sysplex does not necessarily include a 9672, but if it does, you need to plan to run the applications that require these features on a different processor, or convert those applications so that they do not require the features.

Complete the column labelled *Need ICRF or Vector Facility* in Table 6 on page 99 by putting:

Y if the application needs either of these facilities
N otherwise

4.2.5 Total Capacity

In cases where workload is moved to a more powerful processor, the capacity planning and workload management considerations are well understood. Basically, the sum of the application capacity requirements determines the required processor type and model. However, if your workload is to be spread across multiple processors, as in a Parallel Sysplex, you may be facing an unfamiliar scenario. You will have to distribute the workload over multiple systems.

In the column labelled *Total Capacity* in Table 6 on page 99, you should fill in the processor capacity required to run each of your applications. You can do that in any units that you wish, for example the processor utilization on a current processor. You can convert this into utilization on a different processor using IBM′s measurements published in the IBM *Large Systems Performance Reference* (LSPR).

The information collected here will be useful to a hardware planner in selecting the appropriate processor configuration for these applications.

In addition to collecting the processor capacity requirements of each application, you also need to consider:

- The ability of this workload to be split over multiple processors

  Can this workload be divided among, say, two MVS systems, each one providing half the required processor capacity? Is there an uneven workload split requirement such that, for instance, 30% of this workload needs to run on a single MVS system, while the other 70% can be split among systems in an arbitrary way?

- Affinity of this workload to other workloads

  Are there any workloads that need to run on the same MVS system as this workload, perhaps because there is no provision to share some of the data?

You may need to do some research to gather this information. The peak and average capacity required by the total workload can be measured using RMF workload reports. Analyzing SMF data can help to determine processor requirements at the batch job step level. Transaction rates and processor consumption, in total or by transaction type, can be measured using specialized tools for the transaction monitors, such as CMF (CICS Monitoring Facility) for CICS, and DC-Monitor for IMS, and evaluation tools like CICSPARS and IMSPARS.

4.2.6 Uniprocessor Power

When you move workload into a Parallel Sysplex environment, you may be making use of processors that have a smaller uniprocessor power than your current processors. This could happen if, for example, you are moving workload from an IBM 9021-711 based processor to an IBM 9672. In this case you need to consider the uniprocessor power required by each application. For example, a CICS region runs largely on one processor engine. You need to evaluate the processing power needed by that one region, since it has to be accommodated on one processor engine.

In the column labelled *Uniprocessor Power* in Table 6 on page 99, you should fill in the uniprocessor capacity required for each of your applications. Be sure to keep the units consistent with the *Total Capacity* column.

You need to consider what a reduction in the uniprocessor size may have on your workload. The effect is a longer processor time for the work. For most workloads, this will have an insignificant effect on transaction response times and batch job run times. For some processor intensive workloads, it may cause a more significant effect. In this case, you may wish to restrict such workloads to run on one of your processors that has a larger uniprocessor power.

4.2.6.1 Response Time Considerations

For many workloads, the contribution of the processor time to the total response time of transactions, and run time of batch jobs, is not significant. Often the network delay accounts for a large part of transaction response time, and DASD I/O accounts for a large part of both transaction response time and batch job run time.

A typical transaction response time is shown in Figure 7. *P* is the processor time, and *PQ* is queuing time waiting for the processor. Typical durations for each part of the response time are given in milliseconds. The total response time is 880 milliseconds.

```
┌─┬──┬───────────────┬──────────────────────────────┐
│P│PQ│      I/O      │           NETWORK            │  880 msec
└─┴──┴───────────────┴──────────────────────────────┘
20 60     300 msec              500 msec
```

Figure 7. Typical Transaction Response Time

If you move this transaction to a system with sufficient total processor capacity, but with a smaller uniprocessor power, even four times smaller, the processor time *P* will increase, but the time queuing for the processor *PQ* may decrease because the workload is being balanced across multiple processors. The overall transaction response time may not change much. It can still be sub-second, as shown in Figure 8, where the total response time is now 920 milliseconds, an increase of only 5%.

```
┌──────┬───┬─────────────────┬─────────────────────────────┐
│  P   │PQ │      I/O        │           NETWORK           │ 920 msec
└──────┴───┴─────────────────┴─────────────────────────────┘
  80    40      300 msec                 500 msec
```

Figure 8. Typical Transaction Executed on a Slower Processor

Similar considerations apply to batch, where the elapsed time for a batch job can be dominated by processor queue time, I/O time, and print delay.

You need to investigate the effect of running the work on smaller engines. In many cases, you will find that there is no significant response time or elapsed time degradation. If you need help to do this investigation, call your IBM Marketing Representative. They have access to tools that project the effect of a smaller uniprocessor power on your specific workload.

TSO: For TSO, high processor time for a TSO transaction can dictate the use of a fast engine. So these users have to be isolated and should not be moved to a system with a significantly smaller engine size. The amount of processor consumption for a given TSO session can be investigated by looking at the SMF record type 30. Groups of users can be placed in a RMF reporting group to see what their processor capacity needs are and what the effect will be of running on a smaller processor.

Batch: The common batch constraint is the need to get certain batch jobstreams completed within the batch window. More and more customers are already taking steps to *parallelize* their batch jobstream due to the shrinking batch window. The more jobs that run in parallel, the shorter the time that the batch run will take. To parallelize your batch stream, you have to look at the relationships between the jobs: some jobs are dependent on the completion of other jobs. OPC/ESA can handle these relationships and can help to optimize the parallelism of the batch stream, in a multisystem as well as a single-system environment. MVS/Batchpipes can also be used to reduce I/O and run more jobs in parallel, thus reducing the batch elapsed time.

4.2.7 Subsystem or Transaction Manager

In this column of Table 6 on page 99 note the subsystem (for example TSO or batch) or transaction processing manager (for example CICS or IMS/TM) used by each application. This is used in the next chapter.

4.2.8 Database Manager

In this column of Table 6 on page 99 note the database managers (for example IMS/DB, DB2, VSAM) used by each application. This is used in the next chapter.

4.2.9 Other Columns

The remaining columns in Table 6 on page 99 are discussed in the next chapter.

Chapter 5. Application Subsystem Support of Parallel Sysplex

In the previous chapter we considered your applications, and completed an inventory of them using Table 6 on page 99. As part of the inventory, we considered which subsystems, such as database and transaction managers, are used by the applications.

In this chapter we will look at the application subsystems, and discuss whether they will enable you to exploit data sharing and dynamic workload balancing. At the time of writing, application subsystems (for example, CICS, IMS, TSO) support Parallel Sysplex functions to different degrees. Although these systems will evolve over time, knowing the currently available functions of each subsystem will help you put together your migration roadmap.

We will look at these topics:

- Exploitation of data sharing
- Exploitation of dynamic workload balancing
- Prioritizing your applications for Parallel Sysplex exploitation
- Mapping your applications onto your processors

5.1 Exploitation of Data Sharing

This section will help you to decide whether your application database manager will support data sharing. Data sharing was discussed in 2.2.1, "Data Sharing" on page 8, and is a prerequisite to dynamic workload balancing. Data sharing means that database updates can be done from different MVS systems while ensuring the integrity of the data. The coupling facility and the database subsystems that support data sharing provide this function by a locking mechanism, whereby updates are synchronized and buffer validity is maintained when an update occurs.

We will consider each of the database managers in turn. Look at the database manager names that you put in the column *Database Manager* in Table 6 on page 99. Note that some applications may use multiple database managers, for example IMS/DB and VSAM. You need to identify whether each one supports data sharing, and what is the relative mix of the multiple types of database access. Techniques for handling the situation where one database manager can do data sharing and one cannot are discussed in 5.2.1.4, "Handling Affinities with CICS" on page 37 and 5.2.2.2, "Handling Affinities with IMS/ESA TM" on page 39.

After reading this section, you should be able to complete the column labelled *Exploitation of Data Sharing* in Table 6 on page 99 by putting:

Y if the database manager supports data sharing now
F if the database manager will support data sharing in the future
N otherwise, so any data sharing is at the data set level rather than at the record level

5.1.1 Data Sharing with IMS/DB

IMS/ESA Version 5 Database Manager (IMS/DB), also referred to as DL/I, has had data sharing for some time, but it was limited to two MVS images. Now, by using the coupling facility, IMS/DB no longer has this limitation and IMS/DB will support high-performance multisystem data sharing in the IMS/TM or CICS environment.

IMS/DB Version 5 supports data sharing for VSAM, OSAM and most types of IMS Fast Path Data Entry (DEDB) databases. Main Storage Databases (MSDBs) and DEDBs that use Virtual Storage I/O (VSO) or have sequential dependent segments (SDEPs) are not candidates for data sharing at the time of writing.

IMS database managers on different MVS systems can access data at the same time. By using the coupling facility in a Parallel Sysplex, IMS DB provides efficient data sharing for more than two MVS systems. If you don′t have IMS Version 5 yet, you can share two IMS systems now by introducing IRLM. This will give you a good starting point for data sharing with IMS Version 5.

We will briefly discuss the components of IMS/DB data sharing.

IRLM: IRLM (Internal Resource Lock Manager) is the resource locking component of IMS (and DB2). It is an independent component shipped with IMS/ESA DB Version 5. IRLM 2.1 uses the coupling facility′s lock structure for its locking function. Each IMS data sharing group needs at least one IRLM per MVS.

DBCTL: For CICS transactions to be able to share IMS/DB data in a Parallel Sysplex, the CICS system must communicate with the IMS/DB system using the DBCTL interface. DBCTL (Database Control) controls the database I/O in an IMS/DB system. One DBCTL IMS Region per MVS system is needed in a Parallel Sysplex to exploit IMS data sharing. If you have more than one DBCTL in the same MVS image, they can share the same IRLM. If you currently run Local DL/1, you need to migrate to DBCTL. This is primarily a change in CICS definitions.

DBRC: IMS/DB must use Database Recovery Control (DBRC) for database integrity and recovery in a Parallel Sysplex data sharing environment. DBRC is started automatically by DBCTL to control backup, recovery and access for all IMS databases that are registered to DBRC. These functions are essential for data sharing. If your CICS system currently uses another method of database recovery management, you will need to migrate to DBRC. This requires changes to IMS parameters and operational procedures, but does not require any application change.

5.1.1.1 IMS/TM with IMS/DB Data Sharing

Figure 9 on page 29 shows how an IMS/TM system uses IMS/DB to provide data sharing. In the figure, there are three MVS systems and three IMS subsystems that share an IMS/DB database. The IMS Control Regions communicate using multiple systems coupling (MSC) to divide the network over the IMS systems. This also improves the availability in the event that one IMS system should fail. The IMS systems are cloned, so any application can run in any message processing partition (MPP). DBCTL, DBRC and IRLM are needed for database control, backup/recovery and locking.

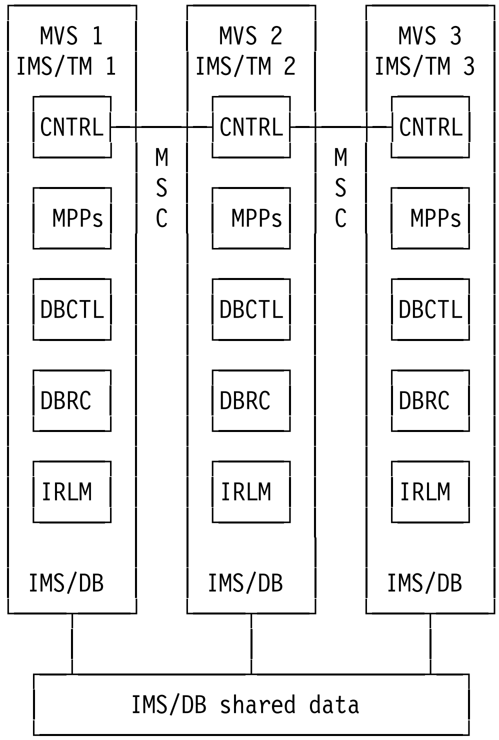

Figure 9. IMS/TM with IMS/DB Data Sharing

5.1.1.2 CICS with IMS/DB Data Sharing

Figure 10 on page 30 shows how a CICS system uses IMS/DB to provide data sharing. The CICS system is split into terminal owning regions (TORs) and application owning regions (AORs). DBCTL is needed, together with IRLM, for locking and DBRC for database recovery functions. The database is shared between all of the MVS systems, so any AOR can reach the data.

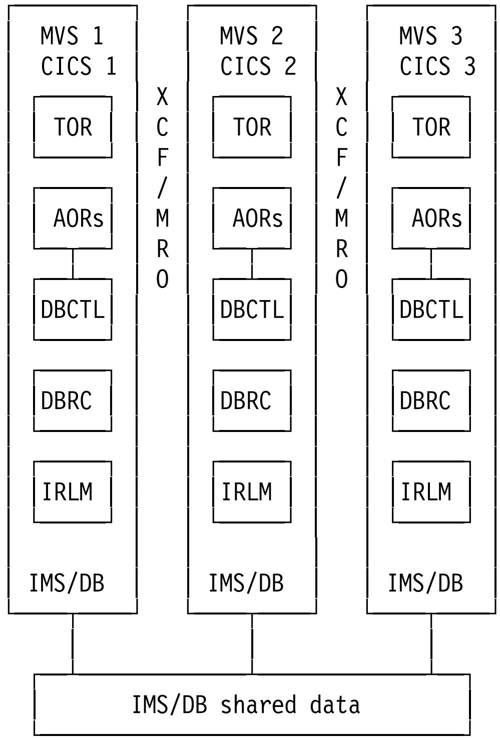

Figure 10. CICS with IMS/DB Data Sharing

5.1.2 Data Sharing with DB2

DB2 Version 4.1 supports data sharing, allowing multiple DB2 subsystems within a Parallel Sysplex to concurrently access and update shared databases. DB2 Version 4.1 became generally available in November 1995.

DB2 data sharing uses the coupling facility to maintain state information (namely status of databases), to lock database records using IRLM, and to buffer (cache) shared data. DB2 uses a store-in cache in the coupling facility′s cache structure, that maintains changed data until it is written to DASD. So if buffers are invalidated because of an update by one DB2 system, and another DB2 system has to read its buffers again, it can just read from the coupling facility instead of reading from DASD. As an option, unchanged buffers may be maintained in the coupling facility as well.

Note that to get to DB2 Version 4, you must migrate from DB2 Version 3. Therefore DB2 Version 2 customers need to migrate to Version 3 first, then to Version 4.

As with IMS/DB, DBRC is used for database control and recovery, and IRLM is used for locking between multiple MVS systems. The DB2 implementation looks similar to the IMS/DB implementation shown in Figure 9 on page 29 and Figure 10.

5.1.3 Data Sharing with VSAM

In September 1994, as part of the announcement of *Version 4 of IBM CICS for MVS/ESA*, (program number 5655-018), IBM made the following statement of general direction:

> To facilitate customer planning, IBM intends to enhance CICS/ESA to provide VSAM data sharing between CICS regions. This will include VSAM record level sharing (RLS) between multiple CICS systems and centralized sysplex-wide log and journal management. Support for the underlying facility will be provided by a future MVS/ESA release. Further Parallel Sysplex support will provide for temporary storage data sharing between CICS systems in the same or different MVS systems. IBM may change its product plans in the future for business or technical reasons.

CICS/VSAM record level sharing is expected to require the VSAM data sets to be system-managed. However, it is only necessary to manage those data sets that are shared.

5.1.4 Data Sharing with Other IBM Access Methods

Workloads that use other access methods, such as sequential (SAM), direct (DAM) and partitioned (PAM), can share data in the conventional way at the data set level (as opposed to at the record level). In conjunction with a multisystem enqueuing mechanism like the GRS component of MVS, this is a well proven way of sharing data for many batch and TSO workloads. Access can be controlled so that, for example, two batch jobs will not update a data set at the same time. This may be an acceptable level of data sharing for some applications, even though you have completed the column labelled *Exploitation of Data Sharing* in Table 6 on page 99 with a *N*.

For the OpenEdition Hierarchical File System (HFS) IBM has made the following statement of general direction:

> Improvements will be made (in a future release of MVS/ESA) to the HFS to allow files to be shared across a sysplex environment. This will improve file availability when using Parallel Sysplex processors.

Note that this statement presents IBM′s currently intended plans. IBM will continue to monitor business conditions and requirements and may make changes to these plans as required.

5.1.5 Data Sharing with Non-IBM Database Subsystems

There is significant interest in providing data sharing capabilities from non-IBM database vendors. At the time of writing, ADAPLEX+, a product of Software AG, supports the Parallel Sysplex and qualifies as a parallel database manager for Parallel Sysplex Licence Charge. You should check the current state of products you are interested in with the vendor.

5.2 Exploitation of Dynamic Workload Balancing

This section will help you to decide whether your application subsystem or transaction processing manager will support dynamic workload balancing. Dynamic workload balancing was discussed in 2.2.2, “Dynamic Workload Balancing” on page 11. You will then be able to complete the column labelled *Exploitation of Dynamic Workload Balancing* in Table 6 on page 99 by putting:

Y if the subsystem supports dynamic workload balancing now
F if the subsystem will support dynamic workload balancing in the future
N otherwise

This is based on the name of the subsystem that you put in the column *Subsystem or Transaction Manager* in the same table.

We will look at:

- CICS
- IMS/TM
- Other subsystems, including TSO and batch

5.2.1 CICS

CICS Version 4 provides support for dynamic workload balancing. In order to implement dynamic workload balancing, you will probably wish to use CICSPlex/System Manager (CICSPlex/SM) as well, as it provides the dynamic routing algorithms and system monitoring facilities that are required for dynamic workload balancing. There are also some application requirements that are necessary for dynamic workload balancing, and we will discuss these in 5.2.1.3, “Application Considerations for Dynamic Workload Balancing” on page 36.

Versions of CICS before Version 4 do not support the dynamic workload balancing provided by CICSPlex/SM. However they do provide many capabilities for splitting terminal, application, and data functions into multiple CICS regions. They also provide a flexible, though static, communication mechanism between them, both within an MVS system and across MVS systems. *All* supported versions of CICS will run in a Parallel Sysplex. However CICS Version 4.1 is the minimum CICS release needed to qualify for Parallel Sysplex Licence Charge.

We will discuss:

- Workload distribution with all versions of CICS
- Dynamic workload balancing with CICS Version 4
- Application considerations for dynamic workload balancing
- Handling affinities with CICS

5.2.1.1 Workload Distribution with All Versions of CICS

CICS provides the facility to split work into different CICS regions. They communicate with each other using transaction routing and function shipping. These techniques are already used in many installations. CICS workload may run in the following types of CICS regions, all of which are standard regions with particular work routed to them:

- Terminal Owning Regions (TORs), dealing with terminal processing.
- Application Owning Regions (AORs), dealing with applications.
- File Owning Regions (FORs), providing access to data such as VSAM.
- Data Base Control Regions (DBCTL) for IMS/DB processing.
- Queue Owning Regions (QORs) that contain queues that are accessible by any AOR. They are defined to the AORs as *remote* queues. This can avoid inter-transaction affinities that might occur with temporary storage or transient data queues.
- Print Owning Regions, dealing with CICS printing in the network.

Transaction routing is used to route transactions from TORs to AORs, and between AORs, and function shipping is used for the communication between the AORs and the File Owning Region. Two different examples of this are shown in Figure 11 on page 33. The first shows all terminals coming into one TOR, and being routed from there. The second example shows the network coming into two different TORs.

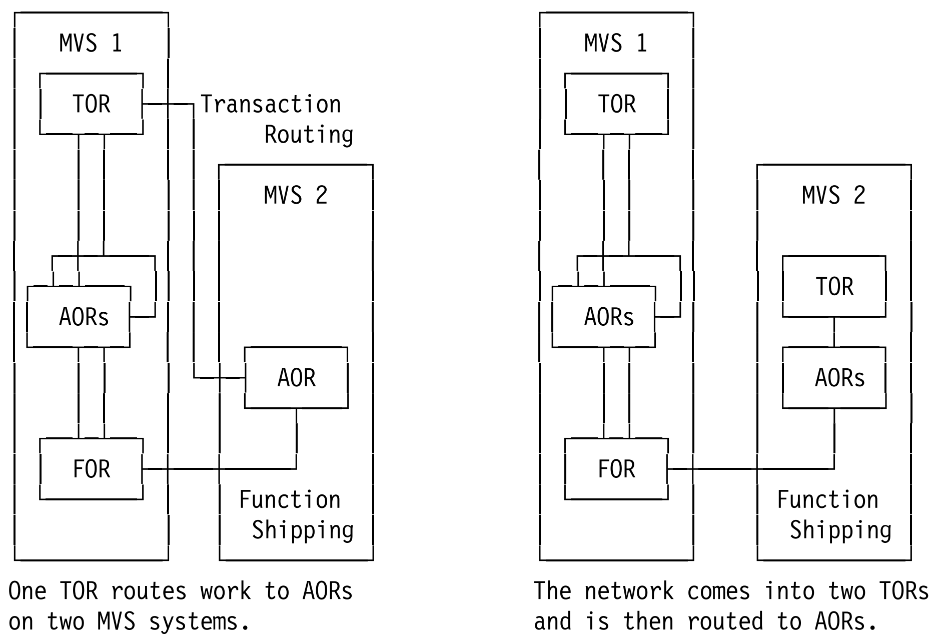

Figure 11. Examples of CICS Workload Distribution

CICS Multiregion Operation (MRO): The communication between CICS regions is achieved using CICS multiregion operation (MRO). With MRO, users can log on to one CICS terminal owning region (TOR), and that TOR will route transactions to various CICS application owning regions (AORs) on multiple MVS systems. This provides an easy user interface, and is the first step towards CICS dynamic transaction routing. MRO is a facility of CICS that has existed for many years, and many customers have already implemented it. The implementation involves changing CICS tables; applications run unchanged.

Communication between CICS regions on different MVS systems (before CICS Version 4) is achieved using Intersystem Communication (ISC), which uses a VTAM connection between the MVS systems. The overhead of communicating between MVS systems this way has precluded some configuration options. CICS/ESA Version 4 introduced cross system MRO, which allows communication between MVS systems using XCF. This option, which is discussed in "Cross-System MRO" on page 35, can be retrofitted to earlier versions of CICS, thus allowing more efficient distribution of CICS workload within a Parallel Sysplex.

Static Transaction Routing: Static transaction routing is available in all current CICS releases. This allows workload to be spread across multiple CICS regions and multiple MVS systems as described above. However the workload distribution is static; that is, it is according to a predefined set of tables. It makes no allowance for which CICS regions are currently available, nor which systems are achieving good performance. It may be suitable for distributing some workloads, and is a first step to exploiting dynamic workload balancing.

So by creating Terminal Owning Regions and Application Owning Regions, the CICS workload can be spread over multiple MVS systems, with or without data sharing and dynamic workload balancing. However, if data sharing is not implemented, the File Owning Region and the DBCTL region cannot be split, and must be placed on a single MVS system.

Spreading CICS Workloads without Dynamic Workload Balancing: This section discusses two examples of spreading CICS workloads using data sharing but not using dynamic workload balancing. These examples are variations on Figure 11 on page 33.

- Example 1

  Suppose you have a single CICS region with several applications, one of which needs a larger uniprocessor power and cannot be split. How can you still offload work to a system with a smaller uniprocessor power and continue to grow?

  Split the smaller applications from the one that needs the larger uniprocessor power. Keep the application with the high engine capacity need on the larger processor in its own AOR. Create AORs for the smaller applications, and run them on the system with the smaller uniprocessor power. Use a FOR for all non-data sharing file access. Files that use data sharing can be accessed directly from database managers running on each processor. The benefits are:

  – Higher availability. With the isolation of applications and the separation of the data from applications, failures will impact fewer applications.
  – More capacity, since a single CICS region is no longer processing all applications and the terminal control functions.
  – Virtual Storage Constraint Relief (if it is still a problem).

- Example 2

  Suppose you have an application that uses VSAM and IMS/DB, and you want to implement data sharing. Data sharing support is available for IMS/DB, but not for VSAM at the time of writing.

  Split the application across multiple AORs on different systems. The IMS/DB databases can be accessed directly with data sharing through a copy of IMS/DB running on each system. CICS MRO function shipping allows all VSAM access to be handled through a File Owning Region (FOR) running on one system. This example is shown in Figure 12 on page 37.

5.2.1.2 Dynamic Workload Balancing with CICS Version 4

CICS/ESA Version 4 introduces cross system MRO and enhances the dynamic transaction routing available in Version 3. We also discuss VTAM generic resource support.

Dynamic Transaction Routing: CICS/ESA Version 4.1 supports dynamic transaction routing. Together with the MVS workload manager and CICSPlex/SM, this provides dynamic workload balancing for CICS transactions. This means that transactions can be routed from any TOR to any AOR if the specific application is available in that AOR.

CICSPlex/SM goes beyond simple workload balancing, which is spreading the work evenly across the available processors. It provides more comprehensive workload management, taking into account affinity considerations and response time goals. CICSPlex/SM selects the AOR that is most likely to meet the response time goals set for the transaction in the MVS workload manager service definition. CICSPlex/SM is alternatively able to select the AOR that has the shortest queue of transactions. For both these methods, CICSPlex/SM maintains information about the state of the candidate AORs, that it obtains by constant monitoring through its agents resident in the AORs. Besides the dynamic workload balancing support, CICSPlex/SM also provides a single point of control for your CICS regions. CICSPlex/SM will work with CICS releases prior to Version 4, without the dynamic workload balancing capability.

Dynamic transaction routing and workload balancing are also the base for continuous availability for CICS applications. If a CICS AOR or an MVS system is not available because of failure or scheduled outage, CICSPlex/SM will note this and route transactions to an AOR that is available to process the transaction.

Cross-System MRO: In a Parallel Sysplex with CICS/ESA 4.1, cross system MRO can be used instead of ISC for communication between CICS regions on different MVS systems. Cross system MRO uses the MVS XCF function, and thereby incurs significantly less overhead than ISC.

The CICS/ESA 4.1 module DFHIRP that offers the function of cross system MRO can be used with earlier CICS releases, thus improving performance of those releases also.

Transaction routing using XCF/MRO can give a 26% to 31% throughput benefit compared with ISC transaction routing across MVS systems. Function shipping using XCF/MRO can give a throughput benefit of between 103% and 126% compared to ISC function shipping across MVS systems. These results are documented in *CICS/ESA 4.1 Performance Comparison With CICS/ESA 3.3* - see B.1.2, “IBM Tools Disks” on page 107 for how to obtain a copy.

VTAM Generic Resource: VTAM Version 4.2 provides generic resource names. A generic resource name is a name by which a group of CICS TORs in a CICSPlex are known to VTAM. This enables terminal end users to log on to a CICSPlex as though it was a single VTAM application. VTAM selects the specific APPLID of one of the CICS TORs that registered as a member of the generic resource, and logs the user onto that TOR. The user needs only to remember one name, and VTAM will log them on to a CICS region that is currently available.

VTAM generic resource also balances the user sessions across the members of a generic resource; that is, across the CICS TORs. This balances the load on the TORs, and will tend to minimize the amount of transaction routing across MVS systems that is needed to balance the workload, thus reducing the performance overhead.

You do not need VTAM 4.2 generic resources if you split your network over the TORs yourself. In this case the mapping of terminal end users to the TORs is static, but dynamic workload balancing will still route the transaction to the appropriate AOR.

5.2.1.3 Application Considerations for Dynamic Workload Balancing

All CICS applications should be easily converted to CICS Version 4, unless they use macro level coding. Macro level coding is only supported up to and including CICS Version 2. However, transaction affinities can constrain dynamic transaction routing. You need to understand the affinities in your applications, and then decide whether to tolerate them, or to eliminate them.

Understand the Affinities: An affinity occurs when a piece of work (for example, a transaction) must run in a particular place in order for it to run successfully. Some examples of affinities are:

- A set of CICS transactions that pass data using CICS temporary storage must all run in the same CICS region, as temporary storage is not currently shared between CICS regions.
- A piece of work that requires a vector or integrated cryptographic (ICRF) feature must run on a processor with that capability.
- A database manager that is not capable of data sharing can only update each of its databases from one system.

The first thing to make clear about affinities is that they are very unlikely to pose a barrier to Parallel Sysplex implementation. Most customers will run with affinities, certainly initially, and perhaps forever. The key point is to understand them so that they can be addressed. For more information see *CICS/ESA Dynamic Transaction Routing in a CICSPlex*, SC33-1012.

Identify Affinities: It is necessary to identify any affinities, such as those listed above. To identify CICS transaction affinities, use the *IBM CICS Transaction Affinities Utility MVS/ESA* (IBM program number 5696-582). For the initial stage of Parallel Sysplex migration, customers do not need to identify *all* of the affinities in their systems. It is sufficient to identify the affinities for workloads that they plan to move into the Parallel Sysplex environment.

Tolerate Affinities, or Eliminate Them: Having identified the affinities, decide whether to tolerate them, or whether to eliminate them. You will probably only eliminate affinities when the benefits outweigh the costs. For example:

- It may be that a small application change would increase parallelism and therefore application throughput.
- The application may have a critical need for higher availability, and eliminating the affinities would allow it to be dynamically routed across several systems.

As applications are rewritten, consider removing affinities at the same time. For all future application development, you should use techniques that avoid affinities. See also the section “Avoiding Inter-transaction Affinity Associated with CICS Queues” in the document *System/390 MVS Sysplex Application Migration*, GC28-1211.

Affinities can be tolerated by ensuring that workloads with affinities only run on a particular system. For example, one customer tackled CICS transaction affinities by putting all transactions with affinities into one CICS region on one MVS system. All other transactions could run in any of several CICS regions on various MVS systems. Clearly this provided all the benefits of Parallel Sysplex for the other transactions. It also provided benefits for the CICS region with the affinities:

- With data sharing for the databases, a failure of the MVS system on which the CICS region is running can immediately be addressed by restarting the CICS region on another MVS system. Database integrity is assured by data sharing, and restart time is much faster.
- Dynamic workload balancing for the rest of the transactions avoided overloading the processor running the affinities.

Even in the longer term, it is likely that many customers will continue to run with some workload affinities. For a small, stable application, they may choose to continue to run it as today, without dynamic workload balancing. It just may not be worth changing. It will work just fine, and the dynamic workload balancing can go on around it.

5.2.1.4 Handling Affinities with CICS

We provide here an example of using multiple CICS application regions with an affinity. The affinity shown is a database manager that is not capable of data sharing. The techniques can be used with other types of affinity also. This example uses CICS multiregion operation (MRO) for function shipping, to route database requests between CICS systems. Thus, it is possible to run just one copy of a nonsharing database manager, and route all database requests to it from other CICS systems. At its simplest, this could look like Figure 12.

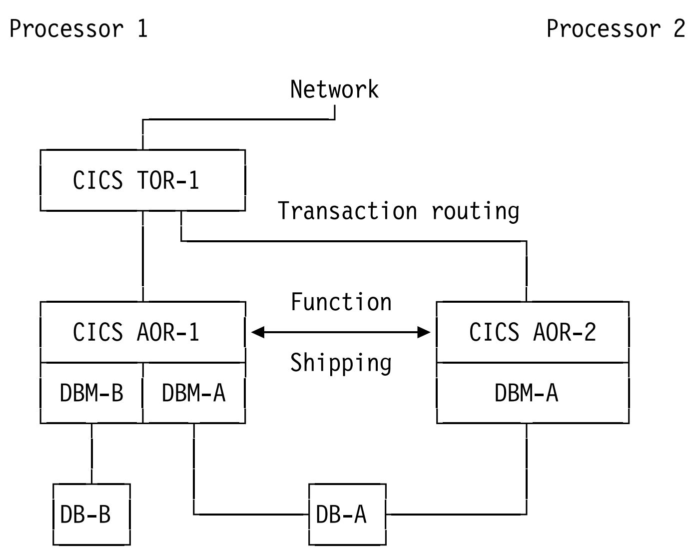

Figure 12. CICS with a Database Manager Not Capable of Data Sharing

Notes:

1. The first database, DB-A, is shared between two copies of the sharing-capable database manager DBM-A.
2. The second database, DB-B, is only accessed by a single copy of the nonsharing database manager DBM-B. Applications running in CICS application owning region AOR-1 can access DB-B directly, while applications running in CICS AOR-2 can access DB-B using CICS function shipping to CICS AOR-1.

3. CICS function shipping uses CICS MRO, and with CICS Version 4.1, MRO can send its messages through XCF that can use the coupling facility. MRO function shipping performance is dramatically improved by using XCF data access rather than ACF/VTAM communication links. The performance results are in *CICS/ESA 4.1 Performance Comparison With CICS/ESA 3.3* - see B.1.2, "IBM Tools Disks" on page 107 for how to obtain a copy. While function shipping would not be used for every transaction, for performance reasons, this does provide a means of splitting workload across multiple CICS and MVS systems, even with a nonsharing database.

4. All terminals could be connected to the terminal owning region CICS TOR-1. This provides a single log on point for all users. Transactions are routed from TOR-1 to one of the application owning regions (AOR-1 and AOR-2); users do not need to know where their transactions run.

5.2.2 IMS/TM

IMS/ESA Version 5 Transaction Manager (IMS/TM) does not currently provide facilities for dynamic workload balancing. IMS/TM does provide facilities for routing transactions between IMS systems, but this is a static routing mechanism.

In April 1994, IBM made the following statement of general direction as part of the announcement of IMS/ESA Version 5:

> IBM intends to further enhance IMS/ESA for the sysplex. These enhancements will enable IMS to further participate in a S/390 parallel transaction environment, with its intrinsic benefits of workload balancing, availability and growth.

5.2.2.1 Workload Distribution with All Versions of IMS/TM

IMS/TM supplies a function called multiple systems coupling (MSC) for communicating between IMS/TM control regions on different MVS systems. One way to use this, without any data sharing, is to split an IMS system into a front-end control region for the applications, running in message processing regions (MPRs), and a back-end control region for the non-data sharing database manager. MSC will route transactions from the front-end control region to the back-end control region. This is a technique that can be used to spread the workload over multiple MVS systems if the database must be kept on one system. It is shown in Figure 13.

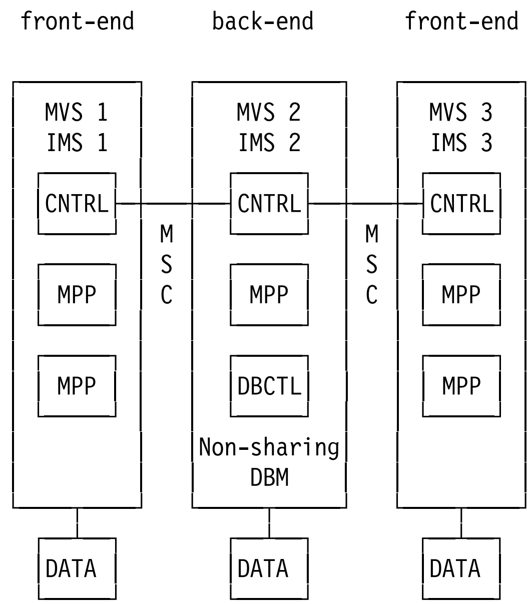

Figure 13. IMS/TM Workload Distribution without Data Sharing

With databases that support data sharing, the IMS workload can be spread across multiple systems in a Parallel Sysplex. You simply clone (replicate) the IMS subsystem and split the network over the cloned IMS subsystems. All IMS subsystems can update the same databases, with data sharing. If one of the IMS subsystems fails, the effect is limited to a subset of the terminals. Note that this is a static workload distribution. The system that your terminal is logged on to is the system where all your transactions will run.

This can be modified using MSC to divert transactions to IMS subsystems on different MVS systems. However this is again a static routing, and requires coding by the installation.

IMS Multisystems Coupling (MSC): It is possible to route transactions between IMS transaction managers using multiple systems coupling (MSC). Users log on to one of the IMS control regions; that control region can run a transaction in one of its own message processing regions (MPR), or it can route it to another IMS control region for running in one of its MPRs. However IMS MSC uses ACF/VTAM links rather than XCF, so customers should use this function sparingly to avoid excessive performance overheads. Also note that multiple system coupling (MSC) is not expected to be the basis for IMS′s future dynamic workload management. See Figure 14 on page 40 for an example of IMS multiple system coupling, used to access nonshareable databases.

5.2.2.2 Handling Affinities with IMS/ESA TM

We provide here an example of using multiple IMS/TM subsystems with an affinity. The affinity shown is a database manager that is not capable of data sharing. The techniques can be used with other types of affinity also. This example uses IMS multiple system coupling (MSC) to route database requests between IMS subsystems. Thus, it is possible to run just one copy of a nonsharing database manager, and route all database requests to it from other IMS systems. At its simplest, this could look like Figure 14 on page 40.

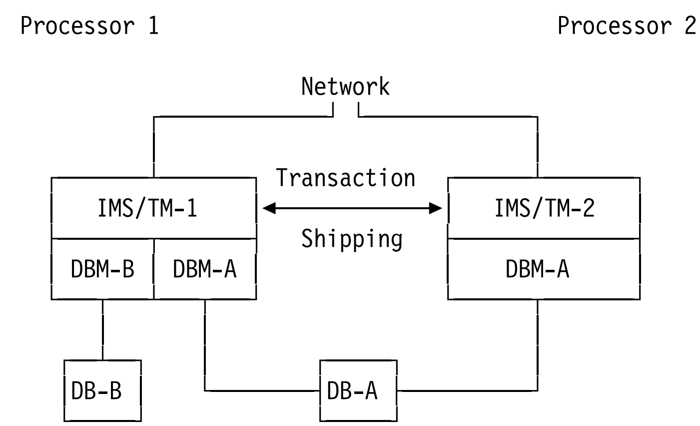

Figure 14. IMS/TM with a Database Manager Not Capable of Data Sharing

Notes:

1. The first database, DB-A, is shared between two copies of the sharing-capable database manager DBM-A.
2. The second database, DB-B, is only accessed by a single copy of the nonsharing database manager DBM-B. Applications running in IMS/TM-1 can access DB-B directly, while applications running in IMS/TM-2 can access DB-B using a customer written transaction that is routed to IMS/TM-1.
3. IMS multiple system coupling (MSC) provides transaction routing; there is no facility analogous to CICS function shipping, though the customer could write transactions to perform a similar function. To use transaction routing the customer must write a routing exit. MSC does not use XCF, so its performance may not be as good as CICS cross-system MRO using XCF.

5.2.3 Other Subsystems, Including TSO and Batch

This section considers other subsystems and transaction processing managers that you may have completed in column *Subsystem or Transaction Manager* of Table 6 on page 99.

5.2.3.1 TSO

TSO does not support dynamic workload balancing. However it is generally fairly easy to split the TSO workload by dividing the users over multiple TSO systems. Users of a specific business unit or department can be grouped together on a particular TSO system. In many cases, every group of users can have their own DASD volumes, so shared DASD will not be required. In other cases, shared DASD can be used to give data set level sharing.

Some TSO users use other subsystems (such as DB2) or applications, and need to be kept on the same systems as those subsystems or applications.

5.2.3.2 Batch Workload Distribution

Batch does not support dynamic workload balancing, although it is possible to route batch jobs to different systems based on job class and OPC/ESA. To do that requires a common spool, such as JES2 multi-access spool (MAS).

Spreading out your batch over multiple MVS systems can introduce the need for shared DASD. GRS can help in properly managing data set level sharing for a batch workload. Data set level sharing is often sufficient for some batch jobs, for example application development.

Note that for IMS/DB batch programs in a multisystem data sharing environment, each batch program requires a separate connection to the coupling facility. However if those programs are run as IMS Batch Message Processing (BMP) programs, all BMP programs running in one IMS subsystem can use one common connection to the coupling facility. So we recommend that you convert IMS batch programs to BMP programs.

5.2.3.3 Non-IBM Subsystems and Transaction Managers

Other subsystems and transaction managers may or may not choose to implement dynamic workload balancing. If they do implement it, then they need to provide intersystem communication to distribute the workload over multiple MVS systems. You need to check the status of your subsystems with the product vendor.

If the subsystem does not support dynamic workload balancing, it may still support data sharing through the database manager it uses, so a static distribution of workloads across multiple systems may be possible.

5.3 Priority - H/M/L/X

Having completed the other columns of Table 6 on page 99, you can now assign a priority for each application to be migrated to exploit Parallel Sysplex. We suggest you use the following priorities:

H if this application has a high priority for exploiting Parallel Sysplex
M if this application has a medium priority
L if this application has a low priority
X if this application cannot exploit Parallel Sysplex

In assessing your priority, you should use the *Business Benefit* and *Risks* columns to assess your priorities, and the other columns to assess the feasibility and amount of work to migrate the application to exploit Parallel Sysplex.

Note that the vast majority of applications are able to exploit Parallel Sysplex without any change. However, if an application is not currently able to exploit data sharing and dynamic workload balancing, then two options are available:

- For a critical business application, the customer may choose to modify it to exploit the Parallel Sysplex environment. This would probably involve allowing greater parallelism, perhaps by removing affinity requirements, perhaps by amending job control statements (JCL). However, the customer would probably only do this if the benefit to the application justified the effort.
- Leave the application alone, and run it in the Parallel Sysplex environment without data sharing and dynamic workload balancing. This application will not get the benefits of those functions, but other applications can. The system level benefits can still be obtained.

5.4 Mapping Applications to Systems

A Parallel Sysplex usually consists of multiple processors and multiple MVS systems. Therefore, at some stage you need to map the applications you identified in Table 6 on page 99 to the different systems. This section discusses the options you have in distributing your workload across multiple systems. In order of increasing function required in the application and application subsystem, the options are:

- Isolating workloads on a single system
- Workload distribution without data sharing
- Workload distribution with data sharing
- Workload distribution with data sharing and dynamic workload balancing

5.4.1 Isolating Workloads On a Single System

The first option on workload distribution is to take an isolated workload and move it onto a separate system, possibly sharing that system with other workloads. By an isolated workload, we mean one that has its own set of address spaces, its own data, and no interdependence with other workloads' address spaces and data. An isolated workload can be moved to a different system, provided that system meets the capacity, response time, and elapsed time requirements of the workload. Figure 15 shows this.

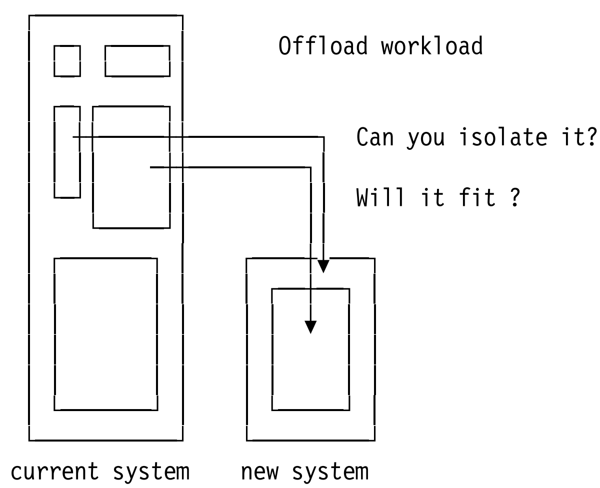

Figure 15. Moving an Isolated Workload to a New System

The capacity of the new system is sufficient if:

- The total processor capacity required by this workload does not exceed the capacity of the new system, and
- The uniprocessor power of the new system is sufficient to meet the response time and elapsed time requirements of the workload.

If the workload is larger than the new system can provide, it may be possible to break the workload into smaller parts that are also isolated workloads. For example, a group of TSO users could perhaps be split into several smaller groups, each of which has their own set of data. Similarly, it may be possible to isolate a subset of CICS transactions and databases from a CICS region.

Isolating a workload can reduce software costs, since it is only necessary to license software on the systems it runs on. Running software on a smaller system can reduce the software cost. It can also simplify device connectivity, by only requiring devices to be accessed by a subset of the processors.

5.4.1.1 Specific Independent Workloads

There are some systems management tasks and system functions that can be treated as isolated workloads. Here is a sample list of discrete workloads that you can offload, for instance, onto a small 9672 within a Parallel Sysplex.

- Network Management

  Several customers have implemented the concept of a CMC machine. This is a separate MVS system that handles all VTAM and network management. This often makes the configuring of the network simpler in a multiple system environment. This function could be isolated, though you need to check that the peak capacity requirements at startup and during recovery can be met.

- Printing.

  You could centralize all printing on one system. If all your systems are connected using a JES2 MAS, then you could run JES/328X on just one system to do all the remote printing. All local printers would be attached to that system, providing simple connectivity. You could also run PSF there.

- LAN management, and LAN serving

  There are a number of tools to help with LAN management. MVS also provides LAN serving capabilities such as LANRES/MVS. Such functions could be isolated to a single system in the Parallel Sysplex.

- Security Management

  All your security administration could be isolated to one system. If you have non-IBM tools that do maintenance and auditing of RACF data, then they can run there. If you use NetView/Access or a similar product, then this can be the gateway into the systems, where all the security checking happens, and then use passtickets to access the other MVS systems. This is discussed further in *Secured Single Signon in a Client/Server Environment*, GG24-4282.

- Systems Management

  Automation can be centralized on one system using Automated Operations Control/MVS (AOC/MVS). Info/Management could run on one system. The controlling part of OPC/ESA can also be isolated, with OPC trackers running on each MVS system. There are a number of systems housekeeping tasks, such as dumps, backups, and processing of SMF data, all of which could be centralized on a single small processor.

5.4.2 Workload Distribution without Data Sharing

It may be possible for you to distribute some workloads that have interdependencies, even without data sharing. For example, a set of TSO users or batch jobs may only need occasional access to common data. This could be handled with data set level sharing managed by GRS. With a CICS workload, transactions can be distributed across multiple CICS AORs, on multiple MVS systems, but can access data through a single CICS File Owning Region (FOR) or a single CICS DBCTL region - see 5.2.1.1, “Workload Distribution with All Versions of CICS” on page 32. Note that this does have performance implications, especially where a lot of shared data is required.

5.4.3 Workload Distribution with Data Sharing

For workloads that support data sharing, their distribution is much simpler. Data sharing is a function of the data base managers, such as DB2, IMS/DB, and CICS/VSAM. With data sharing, a workload can be split up without consideration of data access, since it can get direct access to the data from all systems.

5.4.4 Workload Distribution with Data Sharing and Dynamic Workload Balancing

In all of the previous methods, the distribution of workload is static. That is, a manual adjustment of the distribution mechanisms is required to shift workloads between MVS systems. Dynamic workload balancing provides a way to variably schedule workload on different MVS systems. Workload distribution is dynamic and automatic. The scheduling mechanisms can take into account the current system performance and availability. This allows an application to truly run across multiple MVS systems, with the appearance of a single system image. Dynamic workload balancing requires support from the transaction processing manager.

The session balancing provided by VTAM generic resource can provide some degree of workload balancing across MVS systems, even for transaction processing managers that do not provide workload management at the transaction level.

Chapter 6. Systems Management and Operational Considerations

In this chapter, we will look at some of the systems management aspects of implementing a Parallel Sysplex. The intention is not to provide a comprehensive discussion of all aspects of systems management. Rather, we will focus on things that may change as your installation migrates to Parallel Sysplex. For a detailed discussion of the planning considerations for systems management in the sysplex, including business, change, configuration, operations, performance, and problem management, see *System/390 MVS Sysplex Systems Management*, GC28-1209.

Here, we will focus on two major topics:

- Systems management and operation of a multisystem environment
- Backup and recovery in a Parallel Sysplex

6.1 Systems Management and Operation of a Multisystem Environment

If you have a multisystem environment today, the aspects discussed here will probably be familiar to you. If you have a single-system environment, a Parallel Sysplex will change this into a multisystem environment. Also, if you run only a few MVS images today, Parallel Sysplex, in the long run, may allow you to substantially increase the number of MVS images in your installation. Either way, you need to reconsider the aspects of systems management of a multisystem environment.

In this section, we first look at the traditional systems management tasks that are associated with a multisystem environment. Then we discuss the new aspects and changes that are introduced with Parallel Sysplex.

6.1.1 Traditional Multisystems Considerations

Managing a multisystem environment is a familiar task to some users. The basic approach does not change as you turn your multisystem environment into a Parallel Sysplex. If you have a single system environment today, Parallel Sysplex will change this into a multisystem environment. These are the traditional multisystem issues that you need to consider.

6.1.1.1 Sharing and Switching Peripherals

Sharing means that multiple MVS images can access the same peripheral at the same time. Switching means that access may be given to any MVS image through a switching facility, but only one MVS image has access at a time. You can plan for sharing and switching to be either symmetrical (all MVS images can potentially access a given component) or asymmetrical (only some images can access a given component). Typically you will make your configuration symmetrical at least for switching purposes, while there will be some asymmetry in terms of sharing.

Sharing and switching of hardware components will be necessary and requires planning. For example, a Parallel Sysplex will require you to plan for, install, and share the following hardware components among all MVS systems in the Parallel Sysplex:

- Coupling facilities and the structures that they contain
- Sysplex Timers

- Couple data sets
- Signaling paths between MVS systems
- Other peripherals, especially DASD

Sharing is often restricted by connectivity limitations of certain devices. The following components will typically be switched between all images, but shared by only a subset of images:

- Printers
- Consoles
- LAN and WAN front-end controllers
- Tape devices and robotic tape libraries

Sharing and switching must be managed in an orderly way. If you switch peripherals, you need to direct the workload that accesses these peripherals to the proper system. Even where you share peripherals, you may want to avoid a totally uncontrolled use of certain peripherals by all systems. If a system is taken out of service for an extended time, you need appropriate procedures to switch its peripherals and redirect the workload to another image.

There have been ways to manage sharing and switching in a traditional parallel channel environment for a long time. However many large installations have moved to the ESCON architecture in recent years. A Parallel Sysplex does not require ESCON, but as configurations and switching become more complex, the ESCON architecture provides great advantages in terms of flexibility and manageability.

6.1.1.2 Is Shared DASD Required?

A major advantage of the Parallel Sysplex architecture is its potential to dynamically shift application workload between its MVS images. Therefore, in general, the application data must be available to all images within the Parallel Sysplex. Data base integrity is maintained by the data base management subsystems. For non-database data, like batch and TSO data, data integrity is managed through the Global Resource Serialization (GRS) component of MVS.

Does this mean that all disks must be shared by all systems? No, though symmetrical sharing is easier to manage. Certainly, there will exist applications that can continue to run without the need to share data. These applications may be isolated to their own MVS image within the Parallel Sysplex, if you so choose.

6.1.1.3 Single Point of Control

For a single MVS system, there is a set of consoles to manage that system. As you increase the number of MVS systems, you could simply add consoles and operations personnel for each new image. But clearly, that is not a desirable strategy.

Rather, you want to operate a single system image, with a single point of control. The IBM NetView product provides this capability. It can direct the message and command flow between all MVS images and a central console or set of consoles. You should aim to have one set of consoles for the whole Parallel Sysplex, with traffic to those consoles separated by function rather than by MVS system.

Multisystem Operation: In a multisystem environment, the operators must be trained to handle multiple MVS images. They have to address different systems in operator commands, and they need to control the sharing and switching of peripherals. There are many MVS functions, messages and commands that specifically apply to multisystems operation. Finally, the operator must be aware of and able to control the workload scheduling mechanisms in a multisystem environment.

6.1.1.4 Automation

In large and complex multisystem environments, a human operator is no longer able to reliably manage the complete message flow, and quickly act upon certain events. Automation is required. You can suppress unnecessary messages and block unnecessary alerts, and free operators from console traffic that they do not need. Automation can also react quickly to messages or alerts indicating a problem. This can increase your system and network availability. By replacing difficult or repetitive actions with standard automated procedures, you can reduce operator workload. Finally, by providing installation specific commands and customized panels and displays, you can make the operators′ remaining tasks more manageable.

The IBM NetView product is an ideal platform for your automation strategy. It is supplemented by a number of add-on products, that provide specific automation support for things like:

- MVS console operations
- Network management
- Transaction manager subsystems

6.1.2 New Aspects of Systems Management in a Parallel Sysplex

Although the basic approach to systems management will not change with the introduction of a Parallel Sysplex, there are a number of changes that may affect systems management tasks.

6.1.2.1 Integrating Console Operations

Many customers use the NetView product to establish a single point of control. Starting with MVS Version 4, the Extended Console Services (ECS) function became available as an alternative. It provides a subset of the function of NetView, namely, the ability to direct the console message and command flow of multiple MVS images to and from a single system console. For installations of limited complexity, this is an attractive and cost-free alternative.

However, there are no readily available automation packages for ECS. Therefore, most large and complex configurations will continue to rely on NetView.

6.1.2.2 Centralized Performance Measurement

Certain performance management tasks can be difficult in a multisystem environment. Trying to analyze a performance problem, for instance, may require analysis and comparison of RMF reports from multiple systems.

The RMF component of MVS Version 5.1 provides a function to consolidate performance and workload measurement data within a Parallel Sysplex. RMF data is still individually collected on each MVS image. In addition to writing it out to the SMF data repository, the measurement data for a period of several hours is retained in a processor storage buffer. The RMF post processor program uses XCF facilities to retrieve these buffers from all MVS images in the Parallel Sysplex and consolidates the data in a single listing.

6.1.2.3 Central Control of RACF Status

Managing a multisystem RACF environment is sometimes complicated. Certain control commands, such as switching RACF data sets (RVARY) or controlling the storage buffers within the RACF data space (SETROPT RACLIST or SETROPT REFRESH), must be entered by the operator on each system in the multisystem environment.

In a Parallel Sysplex environment with RACF Version 2.1, the RACF sysplex communication function is available to solve that problem. You need to enter the command only once. The RACF sysplex communications function propagates the command to all other systems in the Parallel Sysplex.

6.1.2.4 Workload Distribution and Balancing

Dynamic workload balancing is done automatically through transaction manager subsystems and MVS services, but you need to define rules to those subsystems that govern these decisions. Maintaining the workload balancing policies is an important new systems management task.

For static workload distribution, the mechanisms that define any desired workload split also need to be maintained. Furthermore, you need to provide the operator with a tested procedure to do this, and you need to document and maintain this procedure.

6.1.2.5 Shared Systems Residence

Maintaining multiple MVS images traditionally results in an increased workload for the systems programmers, as they have to maintain multiple systems residence disks. In a Parallel Sysplex environment, you may choose to maintain only one systems residence disk, and you can share this between any number of MVS images. Customizing is done in a generic way. System dependent parameters in parmlib can be coded as variable symbols, which are resolved into different values by each individual MVS image, based on its system identifier. This is discussed in more detail in 7.1.5, "Software Maintenance Strategy" on page 57.

6.1.2.6 Nondisruptive Systems Maintenance

In traditional systems, maintenance activities are routinely scheduled during nights or on weekends. These activities can be costly. Regular off-shift activities are typically compensated by higher salaries, or by more personnel investment. Furthermore, any software and hardware maintenance activities may decrease the availability of the system. An important task in traditional systems management was to coordinate these maintenance activities, avoiding them where possible, or at least reducing their frequency and duration.

The task remains, but the solution may change in a Parallel Sysplex. If the applications on the system to be maintained allow dynamic workload balancing, you may shift that workload to other systems in a nondisruptive way. You may even do this during normal business hours. If many of these off-shift maintenance activities can be avoided, and if they can be done in a nondisruptive way, it can result in significant savings and increased application availability.

6.1.2.7 Operator Training for Parallel Sysplex

A Parallel Sysplex will require operator training. There are new and modified operating procedures, operator commands, system messages and recovery procedures. The coupling facility hardware, as an example, will introduce the need to sometimes configure coupling facility channel paths online or offline, display information concerning coupling facility resources, as well as shutdown the coupling facility for maintenance activities. MVS/ESA SP Version 5.1 has been updated to support this architecture, including the ability to define the base hardware components and to communicate with the MVS images and subsystems that support them.

6.2 Backup and Recovery in a Parallel Sysplex

The introduction of a Parallel Sysplex can change the way you provide processor recovery. There are new components, such as the Sysplex Timer and the coupling facility, whose recovery have to be considered. Finally, the Parallel Sysplex architecture gives you new options to optimize your approach to the backup in your data center as well as to disaster recovery. In this section, we will look at new backup requirements that are associated with the introduction of a Parallel Sysplex. We consider three topics:

- Backup within a data center
- Disaster recovery
- Combining backup within a data center and disaster recovery

6.2.1 Backup Within a Data Center

Most backup and recovery procedures in a data center, such as those for DASD volumes and communications network components, are not affected by the introduction of a Parallel Sysplex. There are three resources, though, whose recovery aspects have to be looked at:

- Processor recovery
- Recovery of a Sysplex Timer
- Recovery of a coupling facility

6.2.1.1 Processor Recovery

Parallel Sysplex does not inherently require you to provide a backup processor. Due to the high degree of internal redundancy in IBM 9672, 9021 and 9121 processors, processors fail much less frequently today than they used to. If you have no specific processor backup requirement today, you may not need to consider this topic much further.

If a processor fails in a traditional environment, the workload running on it either has to wait until the processor is repaired and restarted, or the workload has to be moved to another processor. In a Parallel Sysplex environment, this changes in three respects:

- Automatic workload takeover
- Backup processor capacity
- Backup processor compatibility

Automatic Workload Takeover: In a pure Parallel Sysplex environment, the Parallel Sysplex architecture can greatly reduce the effect of processor outages, up to the point of completely masking them to the user. In the case of a processor failure, dynamic workload balancing can move workload from the failing processor to the remaining processors. Note that dynamic workload balancing is not implemented in all subsystems yet; see 5.2, "Exploitation of Dynamic Workload Balancing" on page 31.

A new component in MVS Version 5, the Automatic Restart Manager (ARM), can improve this process further. It can monitor the availability of critical applications and subsystems within the Parallel Sysplex, detect a failure immediately, and automatically restart a failing subsystem on the same or on a different processor, based on rules established by the installation.

Backup Processor Capacity: A Parallel Sysplex environment is more likely to consist of machines with vastly different processor capacities than a conventional multisystem environment. Obviously, when planning for processor backup, capacity and performance aspects have to be considered. You need to be sure that if a large processor in your Parallel Sysplex should fail, the workload can either be transferred to another large processor, or that it can be split across multiple smaller processors.

Backup Processor Compatibility: You may run a mixed environment in your data center, consisting of conventional systems (non-coupling capable) as well as Parallel Sysplex (coupling capable) systems. In case of a processor failure, you might wish to move the failing workload to another system. Obviously, the easiest approach is to move the workload to an identical platform. But potentially, this move could take place between conventional and Parallel Sysplex environments. There are two scenarios to move workload to a different environment:

- From conventional systems to Parallel Sysplex

  That is the subject of this whole book. Technically, any workload from conventional systems can run within a Parallel Sysplex, provided the Parallel Sysplex has the required capacity and access to all required resources. But if you have a mixed environment, it may well be because these requirements are not met. You may indeed have some workload that *cannot* run on the currently installed Parallel Sysplex. In this case, you cannot use your Parallel Sysplex as a backup, and you need to provide the backup *within* the conventional systems.

- From Parallel Sysplex to conventional systems

  At the time of writing, any current application subsystem can run on either a Parallel Sysplex or on conventional systems. The new interfaces offered by MVS Version 5, such as XES, are primarily used by the application subsystems. No current application is likely to use such an interface directly; hence applications that run on a Parallel Sysplex can run on conventional systems as well.

  So today, a fallback from a Parallel Sysplex environment to conventional systems is possible. It may require proper adaptation of the application subsystem, for which you have to provide, test, and maintain an appropriate procedure. However, as with any architecture change, you might run into incompatibilities in the future, as subsystems and applications develop and further exploit the new architecture. There may be future application subsystem releases that require a coupling capable processor.

6.2.1.2 Recovery of a Sysplex Timer and a Coupling Facility

A Sysplex Timer can be a single point of failure, so most installations will provide a second timer for backup. If a backup timer exists, and one timer fails, recovery is fully automatic. The backup timer takes over without operator intervention. For more information, see 8.4, "Sysplex Timer" on page 88.

Similarly, a coupling facility can be a single point of failure; therefore most installations will provide a backup coupling facility. If two coupling facilities exist, and one fails, the other one can take over. Each system using the failing coupling facility will copy its version of the contents of the coupling facility structures to the alternative coupling facility. The takeover procedure may require operator action, so it may require some operator training. For more information, see 8.3, "Coupling Facility Selection" on page 82.

6.2.2 Disaster Recovery

A disaster can be defined as a permanent failure of components of a data center, such that no recovery is possible within that machine room. The failing workload has to be taken over by another machine room, probably at a different data center. All data has to be made available (for example, restored) at the backup site and backup processors must be made available to process the workload.

The larger part of disaster recovery planning deals with data recovery, and Parallel Sysplex does not change this. We will focus on processor backup here.

Typically, the processors at the recovery site will not be part of the Parallel Sysplex that was hit by the disaster. Hence, Parallel Sysplex functions will not assist the workload shift to the recovery site in any way. Workload distribution mechanisms only work as long as the participating processors are part of the same Parallel Sysplex, and as long as the data remains intact and accessible.

If failing workloads must be moved from conventional systems to Parallel Sysplex (or vice versa), the same considerations apply as in 6.2.1.1, "Processor Recovery" on page 49.

The Parallel Sysplex architecture will, at least in the medium term future, provide more flexibility in disaster recovery configuration. The workload from a single, large ES/9000 9021 711-based or 9121 511-based processor, for instance, can be moved onto multiple smaller machines after a disaster. Although that move is not automatic, dynamic workload balancing at the recovery site can distribute the workload between the processors available there.

6.2.3 Combining Backup Within a Data Center and Disaster Recovery

A Parallel Sysplex configuration may be planned to provide both backup within a data center and disaster recovery. The basic components of this backup environment are illustrated in Figure 16.

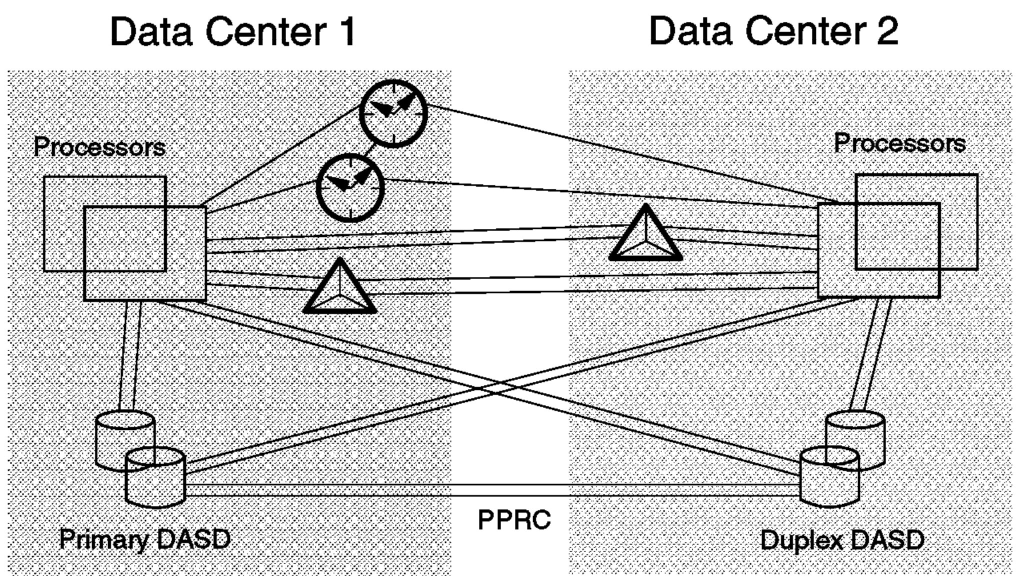

Figure 16. Combining Backup in a Data Center and Disaster Backup. This is just an example. There are many variations of this concept.

Two data centers are placed at a reasonable distance, so that there is little chance that a disaster will strike both at the same time. The distance is limited to 3 km, as this is the maximum supported distance for a sysplex interconnection, and for the distance between a processor and a Sysplex Timer. There are processors in both data centers, configured as a single Parallel Sysplex. There is at least one coupling facility in each data center (note that access to a remote coupling facility will introduce some performance degradation, depending on the distance between the two data centers). You may wish to have at least two coupling facilities in each data center, to provide coupling facility redundancy in a disaster recovery scenario. Two Sysplex Timers are connected together in data center 1, to provide redundancy. Both Sysplex Timers must be in the same data center, since they must be within three meters of each other. All processors have shared access to the primary DASD, which is located entirely in data center 1. The peer-to-peer remote copy (PPRC) feature of the IBM 3990 DASD controller is used to maintain an up-to-date duplex copy of all DASD in data center 2. All processors can access the duplex DASD, if required.

The key advantages of this configuration are:

- The capacity at both data centers can be fully used, since workload can be dynamically scheduled across processors in both data centers using the Parallel Sysplex. Contrast this with many disaster recovery sites, where the processors are often under-utilized (or even completely idle).
- A single component failure in either data center can be more quickly recovered from, again using the Parallel Sysplex features.

For any component failure, a backup component is available to take over the workload. Some of these takeovers are instantaneous, while some require an application restart. In any case, most long outages caused by lengthy data recovery processes can be avoided. In simple terms, this is how it works:

- Recovering from a processor failure

  If any processor fails, dynamic workload balancing can direct the affected workload to other processors in the sysplex. The recovery can be nondisruptive.

- Recovering from a DASD failure

  Failure of duplex DASD in data center 2 is nondisruptive, as it has no immediate effect on any active resource.

  Failure of primary DASD in data center 1 can be recovered by switching to the duplex copy in data center 2. If the application can dynamically reallocate the duplex copy, the recovery is nondisruptive. Otherwise, the address spaces that hold allocations to the failed device need to be restarted to allocate to the duplex device. A small programming enhancement (SPE) is expected to be available soon to allow allocations to be swapped to the secondary device without restarting subsystems. This will use the IOACTION command to suspend I/O activity while the allocation is swapped.

- Recovering from a disaster

  You may choose to run with all the active structures in the coupling facility at data center 1. The coupling facility at data center 2 is therefore used only for backup if the coupling facility in data center 1 fails. In this case, after a disaster at data center 2, the processors at data center 1 can take over the workload dynamically. Primary DASD is not affected, and all lock information is still available in the coupling facility in data center 1, so the takeover can be nondisruptive.

  After a disaster at data center 1, the processors at data center 2 are available to run the workload. They will go into a non-restartable wait state, due to the loss of access to the Sysplex Timers, which are both in data center 1. The lock information for processors in data center 1 is lost (both the processors and the coupling facility are not available). Therefore a full system restart is required at data center 2, including allocating to the duplex devices in data center 2, and backout of inflight transactions at the time of failure. This is the same as if there were no Parallel Sysplex. A backup Sysplex Timer is needed in data center 2 for use in this scenario if there are multiple processors in that data center. Consider setting up alternate parmlib members for running in this situation.

This is obviously just a simplified example. There are many expansions and variations to this concept. In order to work, the concept of resource duplexing needs to include not only processors and DASD, but also things such as network front-ends, consoles, tapes, and printers.

Chapter 7. Software

In this chapter we discuss the software you must implement to run a Parallel Sysplex. Moving to a Parallel Sysplex gives you the opportunity to review how you manage software implementation. It also gives you the opportunity to review what software you have and how you pay for it. We will look at these two topics first.

Then we will look at specific software needed for a Parallel Sysplex. The software shopping list for Parallel Sysplex can appear daunting. However you may only need to install a subset of that software. There are some software products that you need, and others that are optional based on your requirements. We will look at four broad software areas:

- Base software - MVS and related base software. MVS/ESA SP Version 5 is the basic starting point for a Parallel Sysplex.
- Database manager software, including IMS/DB, DB2 and CICS/VSAM.
- Transaction manager software, particularly CICS and IMS.
- Networking and systems management software, such as ACF/VTAM, OPC/ESA, and RACF.

These areas can be considered independently or collectively, depending on your environment. For example, you may choose to migrate to MVS/ESA Version 5 independently of, or at the same time as, you migrate to CICS/ESA Version 4.

7.1 Software Migration Considerations

Parallel Sysplex introduces new functions that can be used in software management. In this section, we discuss changes that you may wish to make in the way you implement software and manage the software environment.

7.1.1 Software Packaging

In some countries IBM has a software package offering that can simplify and reduce the elapsed time of software implementation. You should contact your local IBM representative to establish what is available.

Packaged software can provide the following benefits:

- It is an integrated set of software and hence requires less installation effort, especially where non-IBM software is included.
- It may represent an integrated tested environment and hence reduces your testing requirements.

7.1.2 Non-IBM Software

When reviewing your non-IBM software there are several questions to consider:

- Firstly will it work in a Parallel Sysplex? You may need to install a new level of the software.
- Secondly how do I pay for it in this new environment? See 7.2, “Taking Advantage of Changes in Software Pricing” on page 58.
- Thirdly will it exploit the Parallel Sysplex environment; for example does a database product support data sharing?

As part of its installation planning services for Parallel Sysplex, IBM will assist you with an assessment of your non-IBM software.

If the software doesn′t work, is expensive, or doesn′t exploit the Parallel Sysplex environment, IBM can assist you in evaluating a conversion to an alternative IBM offering.

7.1.3 Minimizing the Migration Time and Effort

A normal migration plan assumes that each component (MVS, CICS and IMS) is tested and implemented separately, which is the way many installations currently implement software changes. However, you can save time and effort by testing all of the new releases together, and implementing them at the same time. The testing time will be longer than for a single product, but there will only be one set of testing rather than three.

This combined implementation would be appropriate when a new production system is being set up, and then workload is being migrated to it, but this is also worth considering when a current system is being upgraded.

7.1.4 Testing the Changes

How do you test the Parallel Sysplex environment? What has changed? Much of the software you implement will need to be tested in the traditional manner. The key difference is the coupling facility and data sharing. The items that need to be tested are:

- Recovery: you need to establish how your recovery procedures will work in the event of a coupling facility failure, a link failure, or the failure of a Sysplex Timer. If you are new to sysplex, then you also need to test the loss of a couple data set. DBCTL may be new to you, and you should familiarize yourself with the recovery of such a system.
- Application changes and affinities: if you change the affinities for an application, then you need to ensure the programs still work and interact correctly. Initially you should plan on removing only those affinities that bring significant benefits.
- Performance: MVS workload manager, data sharing, dynamic workload balancing and new uniprocessor powers will all change the performance characteristics of the system. It is important to stress test your applications to ensure that the performance is satisfactory.
- Operational differences: there will be operational differences. If you get the opportunity to implement a base sysplex early on, then that will give you valuable experience in using a sysplex, in particular the console support, GRS, and couple data set support.
- Disaster recovery: with any significant level of change you should always go back and validate your disaster recovery plan. You should ensure you haven′t introduced changes that will invalidate your disaster recovery plan.

As part of your detailed implementation plan, IBM can assist you in developing a comprehensive test plan to suit your requirements.

7.1.5 Software Maintenance Strategy

Within a single system environment, the requirement to implement software changes is understood. All installations will have a system residence (SYSRES) volume to support IPL, as well as a second volume used to manage software changes. Optionally, a backup of the SYSRES volume may be present for use in the event of a hardware failure.

As you move from a single MVS image to multiple images, the requirement for multiple SYSRES volumes becomes apparent. How should you provide this? Should you simply copy the SYSRES pack, changing values as appropriate to support each new image? When the number of images is relatively small, say four or less, this is a viable approach. But what if your Parallel Sysplex grows beyond four MVS images? Is this still a sound strategy?

Clearly, as the number of MVS images within the Parallel Sysplex grows, your software maintenance strategy must grow as well. You should move away from the strategy of copying SYSRES volumes and move toward a strategy that supports the sharing of SYSRES and the master catalog across some number of MVS images. Support of this strategy requires planning, which we will discuss next.

7.1.5.1 Naming Conventions

In a Parallel Sysplex you will probably be running more systems and subsystems than currently. Naming conventions should be designed with clearly defined objectives in mind. For example, your standard should be clear and meaningful, aid problem determination, allow replication of system definitions, maintain simplicity, support your automation strategy, and allow growth. Does this mean as part of your sysplex migration, you must re-engineer your existing standards? No. It does mean that as additional MVS images and subsystems are added, current standards should be reviewed and altered as it makes sense. One approach to naming conventions can be found in *System/390 MVS Sysplex Application Migration*, GC28-1211.

7.1.5.2 Use of Symbolic Parameters

Symbolic parameters consist of an ampersand followed by a name and can be referenced by system and subsystem components during initialization. Enhancements to MVS provide extensive use of symbolic parameters. MVS/ESA SP Version 5.1 provided support for two symbolic parameters. MVS/ESA SP Version 5.2 provides over one hundred additional symbolic parameters. These additional parameters can be used to support the sharing of not only the various system data sets, but the members within these data sets as well. Doing so will provide for simplified systems management. Use of symbolic parameters is encouraged as you add additional MVS images to the Parallel Sysplex.

7.1.5.3 Sharing the SYSRES Volume

The number of MVS images that can share a SYSRES is limited predominantly by the number of paths to the device. Through the use of ESCON Directors and 3990 model 6s, connectivity becomes less of an issue. Performance is an ever-decreasing concern because I/O to the SYSRES has been significantly reduced due to enhancements to MVS, starting with MVS/XA Version 2. Construction of a shared SYSRES has been simplified as a result of symbolic parameters available in the latest version of MVS.

7.1.5.4 Sharing the Master Catalog

The master catalog is a critical data set that can be shared across MVS images within the Parallel Sysplex when all MVS images are at least at the MVS/ESA SP Version 5.1 level. To aid in the sharing, previously fixed named, non-sharable system data sets can now be specified by the installation using symbolic parameters within the data set name.

7.1.5.5 Availability of Shared SYSRES and Master Catalog Volumes

Before choosing to implement a shared SYSRES or master catalog in your environment, you should consider the implications on availability. If you have only one SYSRES or master catalog volume in the Parallel Sysplex and it fails for whatever reason, then you have lost the entire Parallel Sysplex until you can restart from a backup copy. There are several ways to protect yourself against this:

- If you have more than one copy (cloned copies) of these volumes, and one of them fails, then you have only lost the subset of the Parallel Sysplex that uses that copy. So for example if you have a shared SYSRES for every three systems, then you would lose three systems if one failed. Dynamic workload balancing could dynamically reroute work to the remaining systems in the Parallel Sysplex.
- Another method is to use 3990 dual copy to keep an online, up-to-date copy of these volumes. If one of them fails, the systems will switch to the alternate and continue working.
- A third method is to place the volumes on IBM RAMAC DASD, where RAID-5 operation and dynamic sparing provide protection against hardware failure.

You will probably wish to have at least two SYSRES volumes to enable you to introduce new levels of software into the Parallel Sysplex nondisruptively (one system at a time). The highest availability would come with the use of multiple volumes and RAMAC DASD.

7.2 Taking Advantage of Changes in Software Pricing

With the introduction of Parallel Sysplex, and Parallel Sysplex Licence Charge, there are now several alternative ways of licensing IBM software. We shall examine each of these. There is no simple rule of thumb for predicting the results of the pricing options in your installation. You will need to do a detailed review of the costing of both your IBM and non-IBM software.

Where available, you can now select between the following charge options for each software product:

- Graduated License Charge (GMLC/GOTC) - monthly or one time charge
- Parallel Sysplex License Charge (PSLC)
- Measured Usage License Charge (MULC)
- Indexed Monthly License Charge (IMLC)

You may be able to lower your software costs by:

- Benefiting from Parallel Sysplex Licence Charge, which is an attractive pricing option for software in a Parallel Sysplex. This can reduce the software cost dramatically.
- Reevaluating your software portfolio. You should examine your software portfolio and determine from each vendor how much you will have to pay for

software licenses in a Parallel Sysplex, and on the newer CMOS hardware technology. You should try and choose a software portfolio that minimizes costs while providing the required function. Also, with the trend towards decreasing hardware cost, people costs become more dominant. You should also review your software portfolio to see what software you need to maximize the potential of your people rather than your hardware. You should consider a conversion to an alternative product where appropriate.

Further information on the different pricing options is available from your IBM representative.

7.2.1 Graduated License Charge

Graduated License Charging is the original method of IBM software pricing. Separate charges are assessed for each processor that the product runs on. The charges are based on the software group that each of your processors is assigned to.

7.2.2 Parallel Sysplex Licence Charge (PSLC)

Parallel Sysplex Licence Charge was introduced when Parallel Sysplex was announced, to provide a more equitable form of software pricing across multiple MVS systems. Parallel Sysplex Licence Charge is designed for installations where processors are connected together using a Parallel Sysplex. To qualify for Parallel Sysplex Licence Charge, there are hardware and software requirements that need to be met. The hardware requirements are discussed in 8.1, “Core Hardware Components” on page 77. The software requirements are:

- MVS/ESA Version 5 (JES2, 5655-068 or JES3, 5655-069)
- DFSMS/MVS Version 1 Release 2 (5695-DF1) with full function or any of the alternatives 1 through 5
- ACF/VTAM Version 4 Release 2 for MVS/ESA (5695-117)
- TIOC (5752-VS2) feature numbers 5980/5981/5982 with PTF (these feature codes will supply TIOC only, not the 3.8 base, and are required only by new MVS installations)
- EREP Version 3 Release 5 (5658-260) with PTF
- A parallel transaction manager, one of:
  - CICS/ESA Version 4 Release 1 (5655-018)
  - IMS/ESA Version 5 Release 1 (5695-176) with the IMS TM feature
- A parallel database manager, one of:
  - IMS/ESA Version 5 (5695-176) with the IMS DB feature
  - IBM Database 2 for MVS/ESA Version 4 (5695-DB2)
  - ADAPLEX+ - see the IBM customer announcement letter of December 13, 1994.

As IBM introduces new environments, alternate qualifying products will continue to be defined. Additional non-IBM products may also qualify for Parallel Sysplex Licence Charge in the future.

Parallel Sysplex Licence Charge charges are based on the aggregate capacity, measured in service units, of the processors in the Parallel Sysplex to which the program is licensed. You may have some software that only needs to run on a subset of your processors. For example, you may only have a small CICS system, (with IMS/TM running on the rest of the processors to meet the PSLC qualification criteria). Or you may choose to run some of your systems management software, such as Info/Management, or other software such as compilers, on just one of your processors. By moving these onto a 9672 in the Parallel Sysplex, they can benefit from the smaller capacity, and therefore lower service unit ratings, of some of the 9672 processors.

You may select between GMLC/GOTC and PSLC for a product at the CEC level within a Parallel Sysplex.

7.2.2.1 Parallel Sysplex Licence Charge Extended

In September 1994, IBM announced an extension of System/390 Parallel Sysplex Licence Charge to a single processor. To qualify for Parallel Sysplex Licence Charge Extended, hardware and software requirements must be met. The hardware requirement is a coupling-capable processor running ICMF, as discussed in 8.1, “Core Hardware Components” on page 77. The software requirements are:

- MVS/ESA Version 5 Release 2 (JES2: 5655-068 or JES3: 5655-069)
- DFSMS/MVS Version 1 Release 2 (5695-DF1) with full function or any of the alternatives 1 through 5
- ACF/VTAM Version 4 Release 2 for MVS/ESA (5695-117)

In a single system environment, PSLC charges are based on the individual capacity of the processor. Aggregation of service units is not available in the single system environment.

7.2.2.2 CEC Level Pricing

CEC level pricing is an enhancement to Parallel Sysplex Licence Charge for 9672 E and P models. These machines contain multiple CECs (the 9672 R models contain a single CEC). With CEC level pricing, customers′ charges are based on the capacity of the CECs that the software products use, rather than the capacity of the total machine. Products are licensed to the individual CECs in the 9672. Thus customers can use any combination of CECs to satisfy their workload and at the same time have better control over their software charges. See the IBM customer announcement letter dated September 13, 1994 for more details.

This pricing structure allows customers with 9672 E and P models to:

- Manage their software charges with better granularity
- Add system capacity without affecting software charges for workloads running on a portion of a 9672
- Place new applications on a 9672 at attractive prices
- Place applications on a 9672 and get charges based only on the capacity of the CECs upon which the software is run
- Run versions of MVS other than MVS/ESA Version 5, and other operating systems and the products that run on them, and be charged based on the processor groups of the CECs upon which they run

7.2.3 Measured Usage License Charge (MULC)

Measured Usage License Charge (MULC) is an alternative to Graduated Monthly License Charge (GMLC) or the Parallel Sysplex License Charge (PSLC) for selected IBM S/390 software products. MULC is based on the processor capacity used by a product, measured in CPU service units, which are independent of processor type, size and model group.

MULC applies to a limited set of products. The IBM products currently eligible for MULC are:

- CICS/MVS Version 2 (5665-403)
- CICS/ESA Version 3 (5685-083)
- CICS/ESA Version 4 (5655-018)
- IMS/ESA TM Version 3 (5665-409)
- IMS/ESA DB Version 3 (5665-408)
- IMS/ESA TM Version 4 (5685-013)
- IMS/ESA DB Version 4 (5685-012)
- TSO/E Version 2 (5685-025)

It is IBM′s intention to provide a usage pricing metric for specific versions of DB2.

Eligible products must run under one of the following IBM operating systems to qualify for MULC:

- MVS/ESA SP Version 3 Release 1.3
- MVS/ESA SP Version 4 Release 2
- MVS/ESA SP Version 4 Release 3
- MVS/ESA SP Version 5

MULC can benefit a customer with these products, who:

- Distributes a workload across a Parallel Sysplex
- Runs a small workload on a large machine
- Moves a stable workload to another machine which is in a higher software group
- Consolidates a small workload onto a machine which is in a higher software group
- Places a new application on System/390

7.2.4 Indexed Monthly License Charge (IMLC)

Indexed Monthly License Charge (IMLC) was announced on March 23, 1995 (see IBM Announcement Letter ZA95-0153). It provides a software pricing structure for single processors above 80 million CPU service units per hour (MSUs). Indexed Monthly License Charges will be indexed to the total capacity of the machine on which the software executes, measured in CPU service units. This is consistent with the IBM strategy of pricing software according to value received. PSLC pricing is also based on the CPU service unit capacity of the processor, and IMLC price levels are consistent with PSLC price levels above 80 MSUs.

In order to determine the machines that qualify for IMLC, the maximum capacity of machines placed in processor group 80 is established as 80 million service units per hour (80 MSUs). Processor group 80 includes all machines currently placed in group 80. Software running on a machine of greater capacity will have IMLC prices based on the MSU rating of that machine.

7.3 Base Software

In this section we will look at:

- What are the minimum MVS components you need to implement a Parallel Sysplex?
- Can I do some preparatory work?
- What additional changes should I consider?

By base software, we mean the MVS product and associated software that is considered part of the basic system, such as ACF/VTAM, RACF and JES. We exclude database managers, transaction processing managers, and network and systems management software, all of which are considered later.

7.3.1 The Minimum Required

This section discusses the minimum required to get a Parallel Sysplex working, with XCF and GRS using the coupling facility. XCF (cross-system coupling facility) is the primary communication interface in a sysplex. In a base sysplex, XCF uses channel-to-channel connections (CTCs) to communicate between MVS systems. In a Parallel Sysplex, XCF can use the coupling facility to communicate instead of, or as well as, using CTCs. Global Resource Serialization (GRS), the enqueuing and dequeuing interface in MVS, uses XCF to transport its messages between MVS systems.

There are two tasks in getting to a Parallel Sysplex: migrating to the required software levels, and implementing the Parallel Sysplex.

7.3.1.1 Required Software Levels

To implement a Parallel Sysplex you need the following software levels:

- MVS/ESA SP Version 5.1 (JES requirements are discussed in 7.3.2.3, “Do I Have to Implement the Latest JES?” on page 64).
- DFSMS/MVS 1.2 (there is no requirement to have SMS managed data).
- Related products such as PSF/MVS, RMF, ICKDSF, SMP/E, TIOC, EREP, and others need to be at specific release levels to work with MVS/ESA SP Version 5.1 or later. See the IBM customer announcement letter for the MVS release you are installing for details.
- MVS/ESA SP Version 5.2 requires the High Level Assembler (5696-234).
- The MVS release you are installing may require specific levels of non-IBM products also. You should check with the vendor.

The ordering and installation of the required software levels for MVS and its associated products is the normal process.

Most of the MVS/ESA SP Version 5.1 release changes are added function to enable data sharing and workload balancing. If data sharing is not being used, then a lot of the changes are not noticeable. The main migration consideration for MVS Version 5.1 is the delivery of 4-digit device numbers. Some non-IBM products and user-written code could be sensitive to this change, and need checking.

Many installations will choose to upgrade the whole of the MVS product set at the same time, using one of IBM′s software package offerings. Thus they will also install the new RACF and RMF versions, which will exploit the coupling technology to improve systems management and performance. You may also choose to install the new levels of database manager and transaction manager at the same time; see 7.1.3, “Minimizing the Migration Time and Effort” on page 56.

Notes:

1. The requirement here is to install the correct product levels. It is not necessary to implement all of their function.
2. For a noncoupled processor to qualify for Parallel Sysplex Licence Charge Extended, it must be running MVS/ESA SP Version 5.2.

7.3.1.2 Implementing the Parallel Sysplex

Once you have the software and hardware installed, you can now implement the Parallel Sysplex, using the coupling facility. The majority of the new things you need to do relate to the coupling facility. You will need to establish definitions for the coupling facility and describe how it is to be used. You will also need to set up some couple data sets (which are part of a base sysplex also). These data sets contain status and policy information describing how you want the sysplex to function. You will need to plan for education of systems programmers and operators.

While this work will not be familiar to you, it is not difficult. If you wish, IBM can help you with these steps under a services contract.

7.3.2 Preparatory Work

There are many things you can do with your current version of MVS to prepare for Parallel Sysplex. These depend on which level of MVS you are currently running. For example, you can set up a GRS ring and JES2 multi-access spool (MAS) configuration using versions of MVS prior to MVS/ESA SP Version 5.1.

We will discuss some of the preparatory steps you could take, and also clarify some of the prerequisites for Parallel Sysplex.

7.3.2.1 Do I Need to Migrate To MVS/ESA Version 4 First?

If you are on MVS/ESA Version 3, do you need to implement MVS/ESA Version 4 before MVS/ESA Version 5? Definitely not. The total effort in going from Version 3 to Version 5 will only be slightly more than going from Version 3 to Version 4, and significantly less than going from Version 3 to Version 4 and then to Version 5.

If you are on MVS/XA Version 2, then go directly to MVS/ESA Version 5.

You should review the cumulative effect of the software version changes to ascertain the total changes in processor memory requirements. You should also review any performance changes in the intermediate releases.

7.3.2.2 Base Sysplex

Suppose that you have not installed a base sysplex, as provided in MVS/ESA Version 4. Do you have to implement a base sysplex before you can install a Parallel Sysplex? The simple answer is *no*. Bear in mind though that a Parallel Sysplex is a sysplex.

What would implementing a base sysplex first help with? Operations and systems programming would be the prime beneficiaries in getting used to the change in style of operations. You may be able to implement a sysplex today as long as your processor can run MVS/ESA Version 4 or Version 5. That way you can get early experience of the technology.

You can go straight to a Parallel Sysplex as some installations have done. Naturally you would test it and gain experience with it before putting it into production. The initial use of this form of sysplex might be GRS and MCS consoles.

Implementing a base sysplex will require channel-to-channel connections (CTCs), even on a partitioned processor. This may not be too difficult if you have ESCON channels available. A Parallel Sysplex does not need CTCs because cross system communication can go through the coupling facility instead. Thus if it is going to cost money, cause delay or increase complexity to provide CTCs, then it may be a good idea to go straight to a Parallel Sysplex. For more information on implementing a sysplex, see *MVS/ESA SP V5 Setting Up a Sysplex*, GC28-1449.

7.3.2.3 Do I Have to Implement the Latest JES?

Both JES2 Version 5.1 and JES3 Version 5.1.1 exploit Parallel Sysplex capabilities. A sysplex configuration is *required* for a JES2 multi-access spool (MAS) or a JES3 complex using these levels of software.

However, if you do not want to exploit Parallel Sysplex for JES, you don′t have to implement the latest levels.

The MVS/ESA SP Version 5.1 BCP will run with the JES component shipped with the following earlier versions and releases:

- JES2
  - MVS/SP-JES2 3.1.3 + PTFs
  - MVS/ESA SP-JES2 4.2 + PTFs
  - MVS/ESA SP-JES2 4.3 + PTFs
- JES3
  - MVS/SP-JES3 3.1.2 + PTFs
  - MVS/SP-JES3 3.1.3 + PTFs
  - MVS/ESA SP-JES3 4.2.2 (FMID HJS4421) (shipped after December 27, 1991) + PTFs

You do not have to implement a JES2 multi-access spool (MAS) just because you are implementing a Parallel Sysplex. If you have not implemented a MAS and you are running several JES2 systems, the techniques that you use today will continue to work. In particular you can continue to ship jobs and output around using NJE or data sets.

If you do implement a MAS, then you get the benefits of the shared spool, console support within a sysplex and the shared checkpoint in the coupling facility.

If you decide to migrate to JES2 Version 5, then *all* systems in a MAS must be at least on JES2 Version 5.1. JES2 Version 5.1 can coexist with JES2 Version 5.2, but *not* with any prior releases. A cold start is required to get to JES Version 5.1.

The JES3 global region may be uniprocessor power constrained. You should consider which processor to run this on, to ensure adequate performance.

7.3.2.4 Do I Need the ESCON Manager?

The role of the ESCON Manager in MVS has expanded. It can now be used to control the varying of devices online and offline. This is a key function when the devices have different addresses on the different processors within the Parallel Sysplex. If you are going to develop into such an environment, then the ESCON Manager should be considered.

Depending on your ESCON configuration you may not need the ESCON Manager. If for example you don′t have ESCON directors, then ESCON Manager is not essential. The base level of the product supported with MVS/ESA SP Version 5.1 is ESCON Manager Version 1 Release 2 (5688-008) with PTF; however Parallel Sysplex requires Release 3. The base level of the product supported with MVS/ESA SP Version 5.2 is ESCON Manager Version 1 Release 3.

7.3.2.5 Do I Need to Implement SMS?

DFSMS/MVS Version 1.2 or later is required for Parallel Sysplex, and also for Parallel Sysplex Licence Charge. However, the term DFSMS has many uses, which can cause some confusion. The Parallel Sysplex requirement is met by running the DFSMSdfp component of DFSMS/MVS 1.2; this component is a follow-on to DFP/XA and MVS/DFP. Thus the requirement is for a DFP upgrade to this level.

DFSMS/MVS also contains three optional components:

- DFSMSdss is a follow-on to DFDSS
- DFSMShsm is a follow-on to DFHSM
- DFSMSrmm is new, providing tape management facilities

If you use any of these optional components, they must be maintained at the same level as the DFSMSdfp component.

This does *not* imply a requirement for system-managed data. DFSMS/MVS can handle both system-managed data and non-system-managed data. IMS/DB and DB2 databases do not need to be system managed to be able to share data.

As part of the announcement of CICS/ESA Version 4, IBM stated that it intends to enhance CICS/ESA to provide VSAM record level sharing (RLS) between multiple CICS systems. This is expected to require the VSAM data sets to be system managed, just as they need to be system managed for VSAM compression. However, it is only necessary to manage those data sets that are shared.

There are many good reasons to have data system managed. However, it is not necessary to migrate all the data to be system managed to implement Parallel Sysplex and data sharing. The migration to Parallel Sysplex can proceed independently of the migration to system-managed data.

7.3.2.6 Do I Need to Implement GRS?

Within a sysplex there needs to be a method of ensuring cross system integrity for the enqueue/dequeue process. The IBM preferred method for doing this is with the product Global Resource Serialization (GRS). If you have an alternative product that provides functional equivalence and performance, then you do not need to migrate to GRS. Similarly, MVS/ESA JES2 Version 5.2 provides a tape sharing capability. If you have an alternative product that provides functional equivalence and performance, then you can continue to use it.

If you already have a GRS ring set up, and now need to add a system or two to the ring to accommodate the new system images that will be participating in data sharing for a critical workload, then there are the following considerations:

- MVS can accommodate a GRS ring with Version 3, Version 4 and Version 5 systems in it. The Version 3 systems will use normal CTC communication, and those that are Version 4 or Version 5 will use XCF. Operationally this level of coexistence is not very nice. In MVS Version 4, several functions were added to make multisystem GRS operations better, and dynamic RNL capability was added. Therefore, a more ideal situation is that the non-coupling systems are at least on Version 4. If all of the systems in the ring are Version 4 or Version 5, then it is much easier to set up and manage.
- If a current GRS ring has Version 3 systems in it, you might consider rebalancing the work so that the coupling systems are in one ring, and the Version 3 systems are in another ring. Operationally this could be easier to manage, assuming the workload rebalancing doesn′t complicate things too much.

7.3.3 Additional Changes You Could Make

This section considers tasks that are not required for a Parallel Sysplex, but which you should consider for their benefits.

7.3.3.1 Workload Manager

While it is not a prerequisite to Parallel Sysplex, we recommend that you implement MVS workload manager goal mode once MVS Version 5 is installed. It is a small amount of effort, and will provide better control of systems performance in line with service level objectives. MVS workload manager can provide more effective use of your processor resource even on a single system. An initial implementation may only take one week. Also, goal mode is used in some forms of dynamic workload management, which may be one of the motives for moving to Parallel Sysplex. For more information on workload management, see *MVS/ESA SP V5 Planning: Workload Management*, GC28-1493.

7.3.3.2 Cloning

We discussed the cloning of systems in 7.1.5, “Software Maintenance Strategy” on page 57. If you choose to use cloning of systems to reduce systems support costs, then some work has to be done to set up the master copy of the system and clone other copies. Some tools and new functions are available to support this effort. There are some parameters in parmlib that are system ID dependent, so the appropriate version of the system ID must be inserted into the appropriate copy of parmlib.

7.3.3.3 Additional Exploitation of Parallel Sysplex

Once the hardware and software are in place, implement the Parallel Sysplex using the coupling facility. This requires the implementation of the Sysplex Timer (unless there is only one processor with LPAR partitions in the sysplex), GRS and the couple data sets. The tasks are documented in *System/390 MVS Sysplex Hardware and Software Migration*, GC28-1210. See also *MVS/ESA Sysplex Migration Guide*, GG24-3925. You can then use the coupling facility for MVS functions to provide benefits to the overall system, even before you implement database data sharing.

Figure 18 on page 68 illustrates the MVS exploiters of a Parallel Sysplex (MVS/ESA Version 5). For comparison, Figure 17 on page 67 illustrates the MVS exploiters of a base sysplex (MVS/ESA Version 4). To the left of the figures is the base component and to the right are the users of that component.

	GRS
	MCS Consoles
	DFSMS/MVS
XCF	TSO/E
	VLF
	SLIP
	DAE

Figure 17. Exploiters of Base Sysplex (MVS/ESA Version 4)

Within a sysplex, communication between the MVS images is done using an MVS component called cross-system coupling facility (XCF). In MVS/ESA Version 4, XCF uses CTCs for intersystem communication. Figure 17 shows the users of XCF in MVS/ESA Version 4.

In a Parallel Sysplex the MVS component that manages the coupling facility is called cross system extended services (XES). Figure 18 on page 68 shows the users of XES. As you can see, XCF can now use XES as its means of communications to other MVS systems. In effect XCF now uses the coupling facility to communicate among the MVS images. This means that all existing and new users of XCF benefit from this.

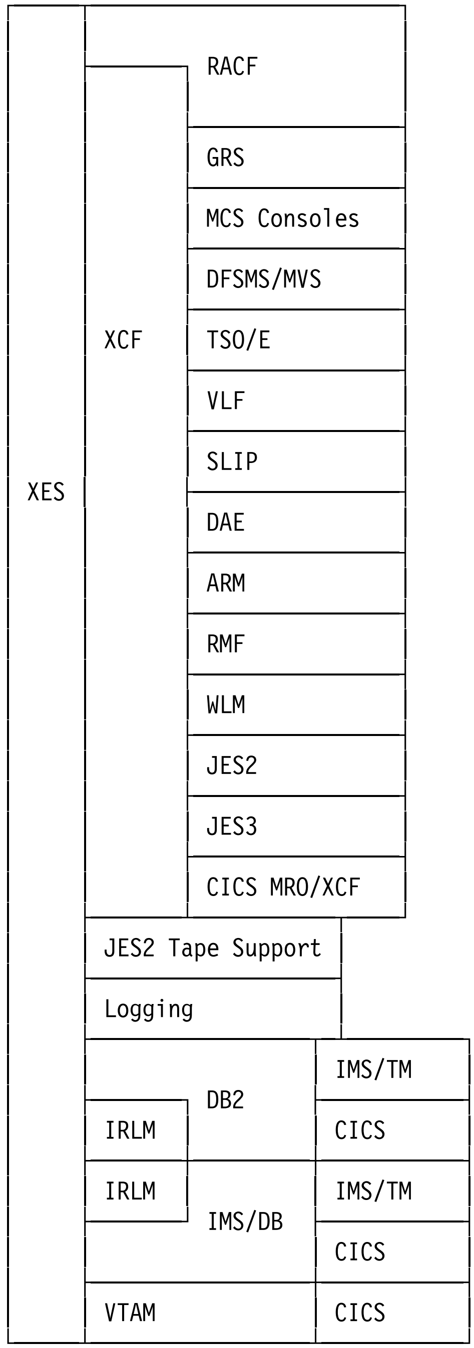

Figure 18. Exploiters of Parallel Sysplex (MVS/ESA Version 5)

Let us look at some of the system level exploiters in more detail to see what levels of software we need, and in what areas they provide benefit.

RACF The RACF control data set can be placed in the coupling facility, providing benefits in performance and systems management. This requires RACF Version 2.1.

JES2 The JES2 checkpoint data set can reside in the coupling facility, providing improved performance. It also provides operational enhancements in the event of a failure. This requires JES2 Version 5.1.

JES3 JES3 inter-system connection (ISC) can use XCF through the coupling facility for availability improvements. This requires JES3 Version 5.1.1.

ACF/VTAM ACF/VTAM uses the coupling facility to keep information about generic resources, enabling multiple systems to look like one to the end user. This enables the single image concept. This requires ACF/VTAM Version 4.2.

RMF RMF can use the coupling facility to provide sysplex-wide performance monitoring. This enables the single image concept. This requires RMF Version 5.1 with MVS/ESA SP Version 5.1, or RMF Version 5.2 with MVS/ESA SP Version 5.2.

CICS MRO CICS/ESA Version 4.1 cross-system multiregion operation (MRO) can use XCF through the coupling facility, providing improved performance.

Tape sharing MVS/ESA SP Version 5.2 supports switchable tape devices that may be shared across systems in a Parallel Sysplex. Automatic tape switching reduces the number of dedicated tape devices required in a Parallel Sysplex and the amount of operator intervention required to manage tape devices.

XCF XCF (cross-system coupling facility) can use the coupling facility, avoiding the complexity of managing channel-to-channel (CTC) definitions with a large number of systems. This is a basic function in MVS. All users of XCF benefit, including:

- Automatic Restart Manager (ARM). This enhances the availability of a system in the event of a system or subsystem failure by automating the recovery actions.
- RACF. This provides automatic synchronization of status and in-storage information across the MVS systems. This requires RACF Version 2.1
- MCS console support, providing a single point of control for MVS software consoles. This requires JES2 Version 4.2.
- GRS. This is a basic function in MVS.
- TSO Broadcast. This requires TSO/E Version 2.4
- OPC/ESA. This requires OPC/ESA Version 1.2

XCF can also continue to use channel-to-channel (CTC) connections rather than a coupling facility.

7.4 Database Manager Software

Data sharing requires the database manager to support it, by using the coupling facility for lock management and buffer integrity across multiple MVS systems. The database manager may need to be upgraded to a new version; this is a “Business as Usual” activity.

7.4.1 Sharing Databases

To achieve database data sharing requires two pieces of work on different processors to be updating the same databases. Section 7.5, “Transaction Manager Software” on page 71 discusses setting up the transaction manager to be able to run one application across multiple processors. However, data sharing can also be achieved with:

- CICS or IMS transactions on one MVS system sharing databases with batch on another MVS system
- Two CICS or IMS applications that are independent from each other except for sharing the database

Having selected an application or applications, establish data sharing for the databases used. Remember that the coupling facility is already in use for MVS functions, so the only new task is to modify database manager definitions.

7.4.2 IMS/ESA Version 5.1 Database Manager

The base levels of the IMS Database Manager that are supported on MVS/ESA SP Version 5.1 are:

- IMS/ESA Database Manager Version 3 Release 1 (5665-408) with PTF
- IMS/ESA Version 4 (5685-012)
- IMS/ESA Version 5.1 (5695-176)

IMS/ESA Version 5 Database Manager (IMS/ESA DB) is one of the qualifying products for Parallel Sysplex Licence Charge (see 7.2.2, “Parallel Sysplex Licence Charge (PSLC)” on page 59).

IMS/DB has been able to share data between two MVS systems for over a decade, using the functions of the IMS Resource Lock Manager (IRLM). Now, IMS/ESA DB Version 5 uses the new IMS Resource Lock Manager (IRLM Version 2.1) to provide greater than 2-way data sharing using the Parallel Sysplex. Parallel Sysplex allows up to 32 systems to concurrently access data. IRLM Version 2.1 is an independent component shipped with IMS/ESA DB Version 5.

The key tasks you are likely to have to do for IMS/DB are:

- Install IMS/ESA Version 5.1 Database Manager
- Implement IRLM
- Implement DBCTL for CICS if not already done
- Implement DBRC if not already used

For CICS with IMS/ESA DB, you will need to migrate to the Database Control (DBCTL) interface if you have not already done so. This is a different interface between CICS and IMS/DB. You also need to implement Database Recovery Control (DBRC) for database recovery and control. These migrations require changes to the product definitions and operating procedures, but do not require any application changes.

We estimate that implementing DBCTL and DBRC, if required, could typically take one month. The installation of a new version of a database manager, including integration with other software, could take three months, and implementing the first shared databases could take a further three months. In all of these activities, the tasks involved are relatively simple and familiar. The key piece of work is testing prior to production running. For more information on this topic, see *System/390 MVS Sysplex Application Migration*, GC28-1211.

Data sharing with IMS/DB was discussed in more detail in 5.1.1, “Data Sharing with IMS/DB” on page 28.

7.4.3 DB2 Version 4

The base levels of DB2 supported with MVS/ESA SP Version 5.1 are:

- Database 2 Version 2 Release 3 (5665-DB2) with PTF
- Database 2 Version 3 (5685-DB2) (with PTF)
- Database 2 Version 4 Release 1 (5695-DB2)

Database 2 (DB2 for MVS/ESA) Version 4 is one of the qualifying products for Parallel Sysplex Licence Charge (see 7.2.2, "Parallel Sysplex Licence Charge (PSLC)" on page 59). It provides data sharing for DB2 databases. Note that to migrate to Version 4, you must currently be running DB2 Version 3. Therefore, you should plan to upgrade to Version 3, if you have not already done so, well before you need to install Version 4.

Data sharing with DB2 was discussed in more detail in 5.1.2, "Data Sharing with DB2" on page 30.

7.4.4 CICS/VSAM

Data sharing with VSAM was discussed in 5.1.3, "Data Sharing with VSAM" on page 31.

7.4.5 Other Database Managers

Many database vendors have stated their intent to provide data sharing for their products using the Parallel Sysplex technology. The delivery of those implementations will be staged over time. At the time of writing, Software AG is the only vendor to have made an announcement of exploitation of Parallel Sysplex by their product ADAPLEX+.

7.4.5.1 ADAPLEX+

ADAPLEX+ is a product of Software AG, and it qualifies as a parallel database manager, extending the Parallel Sysplex to a broader range of environments. See the IBM announcement letter dated December 13, 1994. You should discuss with the supplier what you need to do to support Parallel Sysplex.

While ADAPLEX+ has been certified as a PSLC qualifying parallel database manager, please note that:

- IBM will neither market, sell, nor service ADAPLEX+.
- IBM makes no claims as to the performance, availability, scalability or other characteristics of the ADAPLEX+ implementation based on the Parallel Sysplex technology.

7.5 Transaction Manager Software

As discussed in 7.4.1, "Sharing Databases" on page 69, it is possible to achieve data sharing without transaction manager support, since the data sharing is provided by the database manager. However, transaction manager support is required for:

- Dynamic transaction routing, enabling transactions to be routed based on the current performance of the systems
- Parallel Sysplex Licence Charge qualification

We therefore recommend that you install the correct versions of the transaction managers:

- CICS/ESA Version 4.1
- IMS/ESA TM Version 5.1 (note: this release does not provide dynamic transaction routing).

Upgrading a transaction manager is a “Business as Usual” activity.

7.5.1 CICS

The base levels of CICS that are supported in MVS/ESA SP Version 5.1 are:

- CICS/MVS Version 2 Release 1.2 (5665-403)
- CICS/ESA Version 3 Release 2.1 (5685-083) or CICS/ESA Version 3 Release 3 if using CICSPlex System Manager/ESA (CICSPlex/SM) Version 1 Release 1 (5695-081)
- CICS/ESA Version 4 Release 1.0 (5655-018)

CICS/ESA Version 4 is one of the qualifying products for Parallel Sysplex Licence Charge (see 7.2.2, “Parallel Sysplex Licence Charge (PSLC)” on page 59). Since data sharing is provided by the database manager, and CICS Version 3 runs with IMS/ESA DB Version 5 and DB2 Version 4, CICS Version 3 applications can share data. However, to obtain all of the benefits of Parallel Sysplex, Version 4 is required.

For CICS transactions to be able to share IMS/DB data in a Parallel Sysplex, the CICS system must communicate with the IMS/DB system using the Database Control (DBCTL) interface. This is a new interface introduced with CICS Version 3, as an alternative to the local DL/I interface. You therefore need to migrate your data sharing CICS systems to DBCTL if you have not already done so. This requires changes to CICS parameters and operational procedures, but does not require any application change.

IMS/DB has other requirements that are particularly relevant in a CICS environment, because CICS users may be using alternative ways of managing IMS/DB databases. See 7.4.2, “IMS/ESA Version 5.1 Database Manager” on page 70 for details.

CICS Version 2 systems are able to run in a Parallel Sysplex. However, CICS Version 2 systems are not able to participate in IMS/DB data sharing, because CICS Version 2 does not support the DBCTL interface. CICS Version 2 systems are able to participate in DB2 data sharing. CICS Version 2 does not have a dynamic transaction routing capability. The main reason you may still have CICS Version 2 systems is that it is the last version of CICS that supports macro level programs. However, IBM has now given notice of the discontinuance of program services for CICS Version 2, effective December 31, 1996. You should therefore plan to convert all remaining macro level programs to command level; the benefits of Parallel Sysplex are an added incentive to do this. In the meantime, you should treat CICS Version 2 systems with IMS/DB databases as a specific type of affinity. Databases can either be exclusively owned by one CICS Version 2 system, or can be shared among CICS Version 4 (and Version 3) systems, with function shipping used to gain access from all CICS systems (see Figure 12 on page 37).

All CICS versions are able to benefit from the CICS multiregion operation (MRO) performance enhancements using XCF. The CICS inter-region program (DFHIRP) is a common module used by all CICS systems on one MVS system. Therefore, once the CICS Version 4 copy of this module is installed, all CICS systems will gain the benefits.

For more information on how CICS exploits the Parallel Sysplex to provide dynamic workload balancing, see 5.2.1, “CICS” on page 32.

7.5.1.1 Setting Up Multiple Application Regions

If you want an application to run on more than one MVS system, then it must be able to run in application regions on different MVS systems. You may need to set up additional CICS application owning regions (AORs). Alternatively, you may just need to move some regions you already have to different MVS systems. Communication between the CICS regions is not essential; users can log on to one of the CICS systems and all of their transactions will run only on that one system. However, enabling inter-region communication will provide better exploitation of the Parallel Sysplex.

7.5.1.2 CICSPlex System Manager/ESA

CICSPlex System Manager/ESA (CICSPlex/SM, IBM program number 5695-081) will help in managing a network of CICS systems. These do not need to be in a Parallel Sysplex. CICSPlex/SM is not a prerequisite for Parallel Sysplex. However, CICSPlex/SM provides the easiest way to implement workload management of CICS transactions, in conjunction with the MVS workload manager.

You may have ways of *manually* routing transactions already and may choose to implement CICSPlex/SM at a later stage. Alternatively, it may be a product to implement independently of Parallel Sysplex. CICSPlex/SM can manage CICS systems other than CICS Version 4. As such, it provides a single management point for a diverse set of distributed CICS systems.

The use of Parallel Sysplex may increase the number of CICS regions you run, and in that case the benefit from using CICSPlex/SM will increase.

7.5.2 IMS/ESA Transaction Manager

The base levels of the IMS Transaction Manager that will work with MVS/ESA SP Version 5.1 are:

- IMS/ESA Transaction Manager Version 3 Release 1 (5665-409)
- IMS/ESA Transaction Manager Version 4 (5685-013)
- IMS/ESA Version 5 (5695-176) with the IMS TM feature

IMS/ESA Version 5 Transaction Manager is one of the qualifying products for Parallel Sysplex Licence Charge (see 7.2.2, “Parallel Sysplex Licence Charge (PSLC)” on page 59).

In April 1994, IBM made the following statement of general direction as part of the announcement of IMS/ESA Version 5:

> IBM intends to further enhance IMS/ESA for the sysplex. These enhancements will enable IMS to further participate in a S/390 parallel transaction environment, with its intrinsic benefits of workload balancing, availability and growth.

Meanwhile, it is possible to route transactions between IMS transaction managers using multiple systems coupling (MSC). However IMS MSC uses ACF/VTAM links rather than XCF, so you may wish to use this function sparingly to avoid excessive performance overheads.

Each IMS batch program in a multisystem data sharing environment requires a separate connection to the coupling facility, while all IMS batch message processing (BMP) programs running in one IMS subsystem require only one connection to the coupling facility. So we recommend converting batch IMS programs to BMP programs, although it is not mandatory.

For more information on IMS/TM, see 5.2.2, "IMS/TM" on page 38.

7.5.2.1 Setting Up Multiple Application Regions

If you want an application to run on more than one MVS system, then it must be able to run in application regions on different MVS systems. You may need to set up additional IMS/TM control regions and message processing regions. Alternatively, you may just need to move some regions you already have to different MVS systems. Communication between the IMS regions is not essential; users can log on to one of the IMS systems and all of their transactions will run only on that one system. However, enabling inter-region communication may provide better exploitation of the Parallel Sysplex.

7.5.3 TSO

The base levels of supported and related software in a Parallel Sysplex environment are:

- TSO/E Version 2 Release 4 (5685-025)
- System Display and Search Facility/MVS (SDSF/MVS) Version 1 Release 4 (5665-488) for MVS/ESA SP Version 5.1
- System Display and Search Facility/MVS (SDSF/MVS) Version 1 Release 5 (5665-488) for MVS/ESA SP Version 5.2
- ISPF Version 3 Release 5 (5685-054)
- ISPF/PDF Version 3 Release 5 (5665-402)
- ISPF/PDF for MVS Version 4 Release 1 (5655-042)
- SystemView Enterprise Performance Data Manager/MVS (EPDM) Version 1 Release 1.0 (5695-101)

7.6 Networking and Systems Management Software

In this section we discuss the implementation of networking and systems management software in a Parallel Sysplex.

7.6.1 Networking Software

The minimum levels that will work with MVS/ESA SP Version 5.1 are shown below. This is only the key subset. Refer to the IBM announcement letter for the MVS release you are using for the full list.

- TCP/IP Version 2 Release 2.1 (5735-HAL) for MVS/ESA SP Version 5.1.
- TCP/IP Version 3 Release 1 (5655-HAL) for MVS/ESA SP Version 5.2.
- NetView Version 1 Release 3 MVS/XA (5665-362) with PTFs (in compatibility mode) or NetView Version 1 Release 3 MVS/ESA (5685-152) with PTFs.
- ACF/TCAM Version 2 Release 4 (5735-RC3) with PTFs or ACF/TCAM Version 3 (5665-314) with PTFs.
- ACF/VTAM Version 3 Release 4.1 for MVS/ESA (5685-085) with PTFs or ACF/VTAM Version 4 Release 1 for MVS/ESA (5695-117) with PTFs. You will probably want to install Version 4.2 to get Parallel Sysplex Licence Charge (see 7.2.2, "Parallel Sysplex Licence Charge (PSLC)" on page 59).

7.6.1.1 ACF/VTAM and Generic Resources

ACF/VTAM Version 4.2 for MVS/ESA, or later, is required for Parallel Sysplex Licence Charge. It is also required for VTAM generic resource support. It is possible to implement Parallel Sysplex without generic resources using an earlier release of ACF/VTAM, but in practice you will probably order and install Version 4.2 along with MVS/ESA Version 5. This will avoid another software change in the future, and position yourself for Parallel Sysplex Licence Charge once all of the other components are in place.

Is VTAM generic resource support necessary? It allows the Parallel Sysplex to act like a single image to the user. If you are implementing CICS/ESA Version 4, then you may wish to take advantage of VTAM generic resources. Other transaction managers may not support VTAM generic resources. If users access the system through TCP/IP, then generic resources are not applicable.

For further information on VTAM generic resources see *VTAM Version 4.2 for MVS/ESA, VM/ESA, VSE/ESA Network Implementation Guide*, SC31-6494.

7.6.1.2 ACF/VTAM and APPN

If you implement VTAM generic resources, it requires the Parallel Sysplex processors to be defined as advanced peer-to-peer (APPN) nodes. You do not have to implement APPN throughout the entire network, just the subset that represents the Parallel Sysplex.

Connectivity requirements will grow as each new MVS image is added to your Parallel Sysplex. As such, a more efficient means of both meeting these requirements and providing an infrastructure that will support dynamic workload balancing must be found. ACF/VTAM Version 4.2 support of APPN and generic resource can offer such a solution. Does this mean that your traditional subarea network must be replaced? Absolutely not. Does this mean that both a traditional subarea network as well as APPN can coexist within the same Parallel Sysplex? Yes. This will require some planning, and IBM can assist in this area. Taking advantage of such services will allow your installation to better address the inevitable constraint that traditional front-end processors will place on a Parallel Sysplex as additional network growth is required.

7.6.1.3 VTAM and VTAM CMC

Many users manage their whole network from a single ACF/VTAM, called a Communications Management Configuration (CMC). Such users should be aware that network startup and recovery can be uniprocessor power constrained, as this is single threaded within VTAM. You may run your VTAM CMC as an LPAR on a large processor today for this reason. In that case you need to be careful if you migrate your CMC workload onto a processor with a smaller uniprocessor power.

7.6.1.4 Subarea Changes to Support Parallel Sysplex

If your initial implementation adds no new MVS images, no changes to VTAM are required. If, on the other hand, additional MVS images are added to your current environment, then appropriate changes to VTAMLST must be made. These changes include the creation of system unique VTAM start and configuration members for each new MVS image added to the Parallel Sysplex, as well as the associated major node definitions.

7.6.2 Systems Management Software

The minimum levels that will work with MVS/ESA SP Version 5.1 are shown below. This is only the key subset. Refer to the IBM announcement letter for the MVS release you are using for the full list.

- RACF Version 1 Release 9 (5740-XXH) with PTFs. You do not need to install Version 2, although Version 2 will exploit Parallel Sysplex (see 7.3.3.3, "Additional Exploitation of Parallel Sysplex" on page 66).
- Target System Control Facility (TSCF) Version 1 Release 2 (5688-139).
- AOC/MVS Version 1 Release 3 (5685-151)
- NetView Version 2 Release 4 for MVS/ESA (5685-111).
- Operations Planning and Control/ESA (OPC/ESA) Version 1 Release 2 (5695-007) with PTFs.

Chapter 8. Hardware

In this chapter we identify the hardware required to support a Parallel Sysplex, and the changes you may need to make. First we will look at the core hardware components required in a Parallel Sysplex. Then we discuss how to select an appropriate processor configuration, coupling facilities and Sysplex Timers. Finally we will look at hardware connectivity for I/O devices: DASD, tapes, network and consoles.

8.1 Core Hardware Components

Hardware components are needed to support a Parallel Sysplex. In addition, there are hardware and software requirements to satisfy the criteria for Parallel Sysplex Licence Charge (PSLC). The software requirements were discussed in 7.2.2, “Parallel Sysplex Licence Charge (PSLC)” on page 59.

To qualify for PSLC on a single processor, the processor must be a qualifying processor running an Integrated Coupling Migration Facility (ICMF), or be linked to an external coupling facility. The currently qualified IBM processors are:

- 9672 - all models
- 9121 511-based processors that are at SEC Level 954
- 9021 711-based processors that are at SEC Level 270 or 420

The coupling criteria can also be met by running the Coupling Facility Control Code (CFCC) on a processor that is logically partitioned.

As an added benefit, multiple processors actively coupled in a Parallel Sysplex qualify for lower PSLC pricing, by allowing aggregation of processor capacity (MSUs) across the Parallel Sysplex. To qualify for this capacity aggregation, you must have:

- Two or more qualifying processors connected using a coupling facility. The currently qualified IBM processors are:
  - 9672 - all models
  - 9121 511-based processors
  - 9021 711-based processors
- A coupling facility, which is the coupling facility control code (CFCC) running in an LPAR in any of:
  - 9674 - a dedicated stand-alone coupling facility
  - 9672 - all models
  - 9021 711-based processors

Coupling links are required to connect processors to coupling facilities. Also, at least one Sysplex Timer (IBM 9037) is required when multiple processors are coupled.

Not all processors within a machine room need to be coupling capable. However, only those that are coupling capable can participate in a Parallel Sysplex, and qualify for PSLC pricing.

8.2 Processor Selection

In this section we look at the different processor configurations that you can use to implement a Parallel Sysplex. Then we will discuss how you can select the best option for you, based on your current processors and on your application requirements.

8.2.1 Processor Configuration Options

In 3.1, “Basic Hardware Options” on page 15, we introduced the basic migration options that you have. In this section we will expand on those ideas, and show how all the options can lead to a full Parallel Sysplex environment.

8.2.1.1 Option 1 - Install a Coupling-Capable Processor to Replace an Existing Processor

This option assumes that you have a processor that is not coupling capable. You can replace it with a processor that is coupling capable. The initial migration is a simple processor replacement. Then you can upgrade your software, and establish a test Parallel Sysplex between LPARs using ICMF. This will give you the benefits of Parallel Sysplex Licence Charge. Before going into production, you may connect the processor to an external coupling facility for improved availability. Future capacity growth can be achieved by adding additional processors to the Parallel Sysplex, or by growing the existing processor. An example of this option is shown in Figure 19.

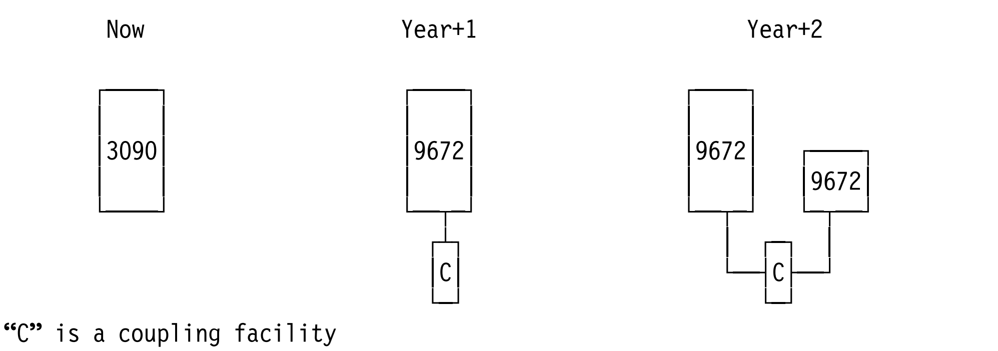

Figure 19. Option 1 - Install a Coupling-Capable Processor to Replace an Existing Processor

8.2.1.2 Option 2 - Install a Parallel Sysplex Alongside an Existing Processor

This option assumes that you have a processor that is not coupling capable. You expand your capacity by installing a Parallel Sysplex alongside your existing processor. The initial migration is to add the additional processors. Workload can be moved onto them, freeing up capacity on the current machine. A Parallel Sysplex is implemented on the new processors, which are eligible for Parallel Sysplex Licence Charge. Future capacity growth can be achieved by adding additional processors to the Parallel Sysplex, or by upgrading processors. Eventually the current processor will no longer be required, and it will be removed. An example of this option is shown in Figure 20 on page 79.

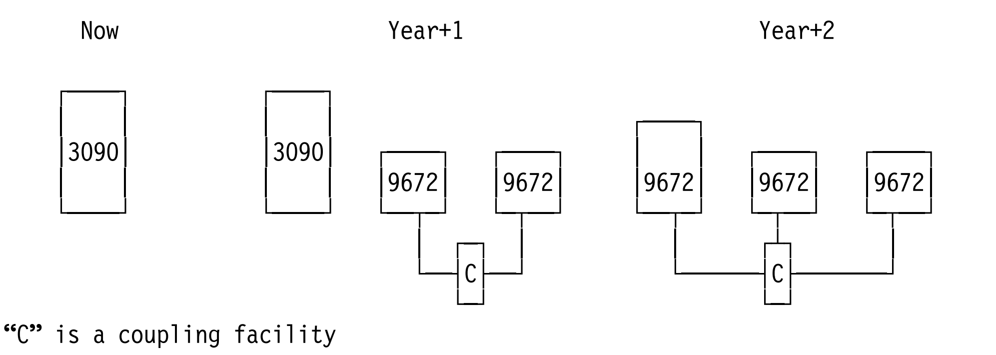

Figure 20. Option 2 - Install a Parallel Sysplex Alongside an Existing Processor

8.2.1.3 Option 3 - Install a New CMOS Processor and Couple It to an Existing Processor

This option assumes that you have a processor that is coupling capable. You expand your capacity by installing additional coupling-capable processors. The initial migration is to add the additional processors. Workload can be moved onto them, freeing up capacity on the current machine. You establish a Parallel Sysplex between the processors. Future capacity growth can be achieved by adding additional processors to the Parallel Sysplex, or by upgrading processors. An example of this option is shown in Figure 21.

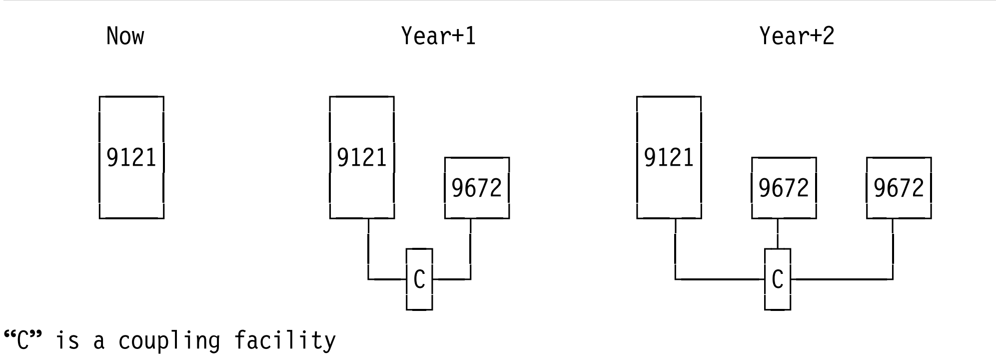

Figure 21. Option 3 - Install a New CMOS Processor and Couple It to an Existing Processor

8.2.1.4 Option 4 - Couple Existing Processors

This option assumes that you have two or more coupling-capable processors, or processors that can be upgraded to be coupling capable. These could for example be 9021 711-based processors or 9121 511-based processors. You add coupling links, and connect them together in a Parallel Sysplex. Future capacity growth can be achieved by adding additional processors to the Parallel Sysplex, or by upgrading processors. An example of this option is shown in Figure 22 on page 80.

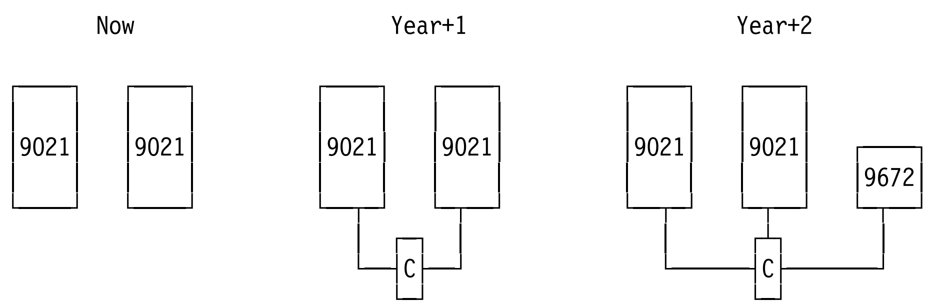

"C" is a coupling facility

Figure 22. Option 4 - Couple Existing Processors

8.2.2 Processor Selection Based on Current Processors

For additional coupling capable capacity you should aim to install the IBM 9672 CMOS processors. These provide the best price performance, and lowest running costs, as discussed in 2.2.4, "Cost Effective Hardware" on page 13. They also provide a finer granularity of upgrade. You may need to take into account the capabilities of the 9672 relative to your current processors, and we will discuss this in 8.2.3, "Processor Selection Based on Application Requirements."

Assuming that your selected applications can run on any coupling-capable processor, you then need to consider:

- Are your current processors able to be upgraded to be coupling capable? If so, then options 3 and 4 may be appropriate choices.
- What financial arrangements do you have for your current processor? If they preclude you from a change, then an alongside Parallel Sysplex as in option 2 would be appropriate. If they invite you to change, then a replacement as in option 1 would be appropriate.
- What is your objective?
  - If you wish to develop a proof of concept model, then a minimally configured test environment may be sufficient. In this case consider small LPARs within an existing coupling-capable processor as in option 3, or within an alongside Parallel Sysplex as in option 2.
  - If you plan to put the selected application into production data sharing as soon as possible, then a fully configured system that addresses single points of failure is required.
- What are the costs of the proposed solution, in terms of hardware, software and people effort?

8.2.3 Processor Selection Based on Application Requirements

After reading Chapter 4, "Applications" on page 19, you may wish to:

- Explore data sharing to prove the concept. If so, you will probably want to keep your initial hardware expenditures low.
- Migrate an existing application to exploit Parallel Sysplex, using data sharing and dynamic workload balancing. If so, you will probably want to configure for maximum availability as the application goes into production.
- Obtain some of the system wide benefits of Parallel Sysplex.

As we have already discussed, to implement Parallel Sysplex you must provide a coupling capable platform. Application requirements will affect which of the four hardware options are relevant to you. You may have completed your application requirements on a copy of Table 6 on page 99. We will use the columns of that table in our discussion here.

Applications requirements may dictate the types of processor that you need in your configuration:

- If some of your applications require a vector facility or an Integrated Cryptographic Feature (ICRF), then you must provide those features on at least some of your processors, and route work that needs those facilities to those processors. See 4.2.4, “Need ICRF or Vector Facility” on page 24 for more details.
- You need to look at the total capacity for each of your applications, to see how you could run it in a multisystem environment (see 4.2.5, “Total Capacity” on page 24).

  If an application is not able to be split using data sharing (see 5.1, “Exploitation of Data Sharing” on page 27), then you may need to provide a processor in the Parallel Sysplex that is at least large enough to run that one application.

  If an application can be split using data sharing, then it can be distributed across multiple processors in the Parallel Sysplex. This would provide benefits as identified in Table 4 on page 97, especially if dynamic workload balancing (see 5.2, “Exploitation of Dynamic Workload Balancing” on page 31) is also supported. You may therefore wish to spread such an application, and you will want at least two separate processors to provide maximum application availability.

  If an application is larger than the largest processor you can obtain, then you will need to find a way to split it across multiple processors, preferably using data sharing and dynamic workload balancing.

  You need to be aware that implementing data sharing and dynamic workload balancing for an application will increase its processor capacity needs. You need to include the cost of data sharing and dynamic workload balancing in your processor capacity calculations. See *System/390 MVS Parallel Sysplex Performance*, GG24-4356 for information on calculating the capacity needed. IBM personnel may also refer to the *MFTSPERF PACKAGE* on MKTTOOLS.
- You need to look at the uniprocessor power for each of your applications, to see the minimum uniprocessor power of the processor it needs to run on (see 4.2.6, “Uniprocessor Power” on page 25). You must provide at least one processor with that uniprocessor power somewhere in the Parallel Sysplex, and route that work to those processors. Note that all processors do not have to have that uniprocessor power, and you may have processors with differing uniprocessor powers in your Parallel Sysplex. You may wish to keep a current processor to handle work with a larger uniprocessor power requirement, and run other applications on CMOS processors that may have a smaller uniprocessor power. This would direct you towards options 3 and 4.

We saw in 8.2.2, “Processor Selection Based on Current Processors” that for additional capacity you should aim to install the IBM 9672 CMOS processors. Application requirements may direct you towards keeping some 9021 and 9121 processors in your Parallel Sysplex.

In looking at the combination of processors in your Parallel Sysplex, remember also to consider the backup configuration in case one of the processors should fail.

We discussed some ideas on mapping your applications onto your processors in 5.4, "Mapping Applications to Systems" on page 42.

8.3 Coupling Facility Selection

You need to provide a coupling facility in your Parallel Sysplex. This section discusses the alternatives for configuring a coupling facility. If you choose to implement a coupling facility within a processor, you will need to add the capacity needs of the coupling facility to the capacity needs for the workloads on that processor.

The coupling facility provides shared storage and shared storage management functions for the Parallel Sysplex. These include locking, high speed caching, and list processing. Subsystems running on multiple MVS systems in the Parallel Sysplex manipulate these shared structures and manage their content through interactions with the coupling facility. This provides an efficient means of sharing data across multiple systems.

A coupling facility is microcode that can run in a separate, dedicated machine, or can run in an LPAR on a S/390 processor. The next section discusses these different options.

8.3.1 Coupling Facility Options

Here we will look at the various ways of implementing a coupling facility, and look at the reasons why you might select each option.

8.3.1.1 Stand-Alone 9674 Coupling Facility

The System/390 9674 Coupling Facility is a stand-alone coupling facility. It is built with the same technology as the IBM 9672. The 9674 provides one processor that is designed to support one or more Parallel Sysplex coupling facilities. The 9674 supports no operating system. It supports only coupling facility LPARs, running coupling facility control code (CFCC), which is a component of its licensed internal code. As with the 9672, the 9674 has a range of processor, storage, and coupling link configurations.

We show an example of a 9674 stand-alone coupling facility in Figure 23 on page 83. This example shows multiple LPARs on a single processor accessing an external 9674 coupling facility. Note that since there is only one processor in the Parallel Sysplex, the partitions can share the same time source without a Sysplex Timer (see 8.4, "Sysplex Timer" on page 88).

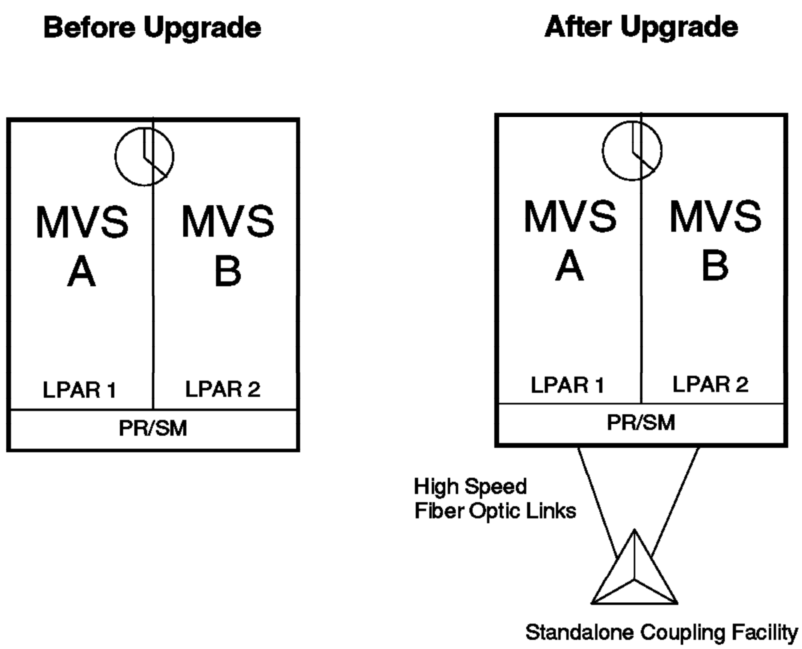

Figure 23. Stand-Alone 9674 Coupling Facility

Why You Might Select a 9674 As a Coupling Facility: A 9674 provides a coupling facility that is external to all processors in the Parallel Sysplex. This provides the highest availability. If any of the processors fail, the information held in the coupling facility is still available to the other processors, and the rest of the Parallel Sysplex can continue to run. If, instead, the coupling facility was an LPAR in one of the processors, failure of that processor would stop the whole Parallel Sysplex, since both copies of information for that processor (within the MVS system and in the coupling facility) would be lost.

A 9674 may also provide the best performance, since the coupling facility is not sharing a processor with other workloads.

The 9674 can also provide the lowest software cost, since Parallel Sysplex Licence Charge does not include the 9674 in the aggregation of capacity in the Parallel Sysplex (see 7.2.2, “Parallel Sysplex Licence Charge (PSLC)” on page 59). If, instead, the coupling facility runs in an LPAR on a processor, the total capacity of that processor is included in the capacity aggregation, even though some of the capacity is used for the coupling facility function.

A stand-alone 9674 coupling facility is therefore likely to be attractive in a production environment, where availability and performance are critical, and the coupling facility capacity used may be significant.

8.3.1.2 Coupling Facility in an LPAR on a 9672 or 9021 711-Based Processor

The Coupling Facility Control Code (CFCC) can run in a logical partition (LPAR) on an IBM 9672 or a 9021 711-based processor. This provides the capability for a logical partition to act as a coupling facility. The coupling facility logical partition is defined in much the same way as today's logical partitions, with the exception that no operating system is IPLed into the coupling facility LPAR. When a coupling facility LPAR is activated, the CFCC is automatically loaded into the coupling facility partition from the processor controller.

Unlike the ICMF option, where coupling links are simulated, a CFCC LPAR requires coupling link hardware to support communication with MVS images running in separate logical partitions on either the same or different processors.

Within the 9672 and 9021, a range of processor, storage, and coupling link configurations are available.

We show an example of LPARs used as coupling facilities in Figure 24. This example shows two LPARs on different processors being used as coupling facilities.

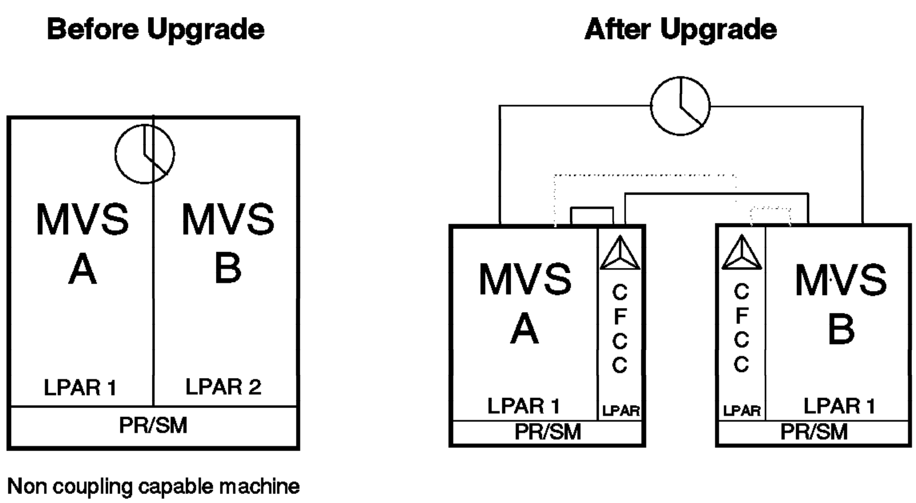

Figure 24. Coupling Facility in LPARs on a 9672 or 9021 711-Based Processor

Note that since there are two separate processors in the Parallel Sysplex, a Sysplex Timer is required as an external time source to synchronize the time of day clocks between the two MVS images (see 8.4, "Sysplex Timer" on page 88).

Why You Might Select an LPAR As a Coupling Facility: An LPAR provides a lower cost coupling facility than using a 9674. This may be appropriate for the early stages of a Parallel Sysplex implementation, where proof of concept and testing are the main requirements. A stand-alone coupling facility could be added later when availability and performance become critical.

You could choose to use a 9672 as a stand-alone coupling facility, with CFCC running in an LPAR, instead of a 9674. In some cases this could be useful when you look at your longer term processor requirements. Note that a 9672 running only CFCC, and no other operating system, has no software license charges.

If you need a large coupling facility, you could choose to use an LPAR on a 9021, though using multiple 9674s may be more cost effective.

In a production environment, you will probably want at least two coupling facilities for availability (see 8.3.4, "Addressing Single Points of Failure" on page 87). As we have seen, for availability the primary coupling facility should be a stand-alone one. The backup coupling facility could be an LPAR on a 9672 or 9021, which is only used if the primary fails. This provides a good combination of availability and cost effectiveness; although note that the LPAR coupling facility does require physical coupling links and adequate processor resource, even when inactive.

8.3.1.3 Integrated Coupling Migration Facility (ICMF)

ICMF is available on the IBM 9672, 9121 511-based processor and 9021 711-based processor. ICMF runs in a PR/SM LPAR, with PR/SM emulating physical coupling links to allow communication between MVS/ESA SP Version 5 systems. When using ICMF, the logical partitions used for the MVS systems and the coupling facility are defined as a group assigned to ICMF. Since ICMF simulates the coupling links and no physical coupling links are used, all logical partitions assigned to the ICMF group must be on the same processor.

Operationally, ICMF logical partitions look like coupling facility logical partitions. All of the same commands and restrictions apply to ICMF as to coupling facility LPARs. A maximum of two ICMF LPARs can be active at once. Use of the ICMF on an ES/9000 requires that the facility be enabled prior to doing a power on reset of the processor in LPAR mode.

We show an example of the use of ICMF in Figure 25. This example shows two LPARs on a processor using ICMF as a coupling facility.

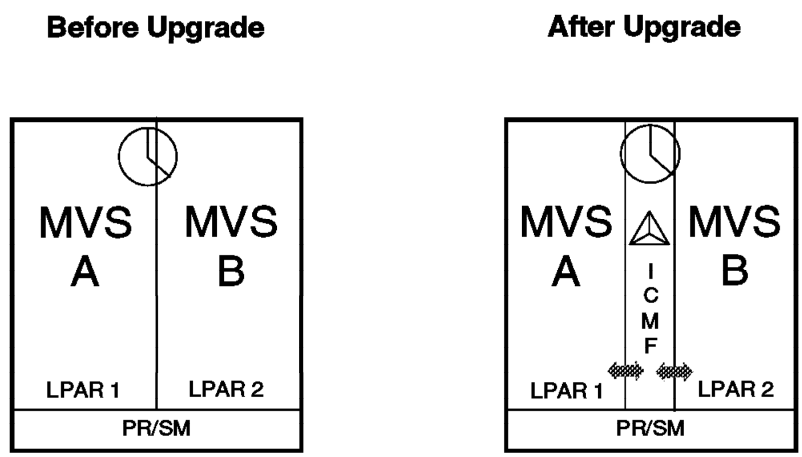

Figure 25. ICMF As a Coupling Facility

All MVS/ESA SP Version 5 systems running on this single processor can share data and signal each other using structures maintained within the simulated coupling facility (ICMF).

Why You Might Select ICMF Instead of a Coupling Facility: ICMF is a very attractive means of supporting proof of concept and testing of a Parallel Sysplex. ICMF simulates coupling facility hardware, so no external coupling facility or coupling links are required. This is therefore a low cost way to implement a Parallel Sysplex. However it does not provide the same level of availability as other options.

ICMF can be used to provide a Parallel Sysplex when you only have one processor. ICMF may also be a useful option in a backup or disaster recovery situation.

8.3.2 Coupling Facility Sizing

You will need to calculate how big a coupling facility you require. The areas you need to consider are:

- Coupling facility processing capacity

  The amount of activity to the coupling facility will dictate the total processing capacity required. Locking and data buffering are likely to be the two major activities. These can vary significantly between different applications and different database subsystems. Each structure must be located in a single coupling facility, but different structures can be placed in different coupling facilities.

- Coupling facility storage size

  Storage must be provided in the coupling facility to hold the structures. You will need to calculate the storage required. Actual coupling facility storage selection must be in the storage increments available on the coupling facility.

- Number and type of coupling facility links

  Transaction locking profiles, buffer hit ratios, the number of MVS images within the Parallel Sysplex, anticipated XCF/MRO traffic, and other factors, will dictate both the type and quantity of coupling links between each MVS image and the coupling facility. We discuss the different types of coupling links in 8.3.3, "Coupling Links."

There are no simple rules of thumb for configuring the size of the coupling facility and the number of coupling links. They are related to the amount of data sharing you are doing. You also need to consider extra capacity to cater for backup if one coupling facility should fail. IBM can help you configure your coupling facility to match your application needs.

8.3.3 Coupling Links

Coupling links provide connections between the MVS systems and the coupling facilities. They are needed for all configurations except ICMF, which simulates the links. The links are high bandwidth fiber optic links.

The links attach to channels on the processors and the coupling facilities. There are two types of links:

- Coupling facility sender links are used by MVS/ESA SP Version 5 images to *send* requests to a coupling facility.
- Coupling facility receiver links are used by the coupling facility to *receive* requests from MVS/ESA SP Version 5 images and process them.

Sender links can be shared between MVS/ESA SP Version 5 images running in LPARs within the same processor. Receiver links must be dedicated to one coupling facility LPAR.

The coupling facility links use a protocol that is Fiber Channel Standard (FCS) Level 0 compliant. This is not the same protocol as used by ESCON. Two types of fiber are available to support these links:

- Multi-mode fiber coupling facility links use either 50/125 or 62.5/125 micron multi-mode fiber, which will transfer information at 50 MBps (million bytes per second).

The maximum distance with multi-mode fiber is one kilometer. Multi-mode fiber is effective for test and emerging production workloads. If the workload manipulates small lock and list structures, such as IMS/DB, multi-mode fiber may be sufficient.

- Mono-mode fiber (sometimes called single-mode) coupling facility links use 9/125 micron mono-mode fiber, which will transfer information at 100 MBps (million bytes per second).

  Mono-mode fiber is required when distance beyond one kilometer is needed or very fast transfer rates are required. The maximum distance with mono-mode fiber is three kilometers. If the application is moving large data structures into and out of the coupling facility, such as DB2 when performing a castout, mono-mode fiber should be used.

It is possible to upgrade from multi-mode to mono-mode fiber, though doing so may impact availability if redundant paths are not available. Unless multi-mode links are needed because of existing cable installations or public access reasons, mono-mode links are recommended.

8.3.4 Addressing Single Points of Failure

In a Parallel Sysplex you must have at least one coupling facility. As with any component, the coupling facility could fail, or it may need to be removed for maintenance activities. For maximum availability, you should therefore configure at least two coupling facilities.

A subsystem builds a structure in one coupling facility. If that coupling facility should fail, the subsystem is notified and will attempt to rebuild the structure in another coupling facility. All subsystems on all MVS systems in the Parallel Sysplex will rebuild the structure from information kept in processor storage. Thus, to achieve high availability, backup coupling facility capability has to be available. The backup coupling facility must have processor and storage resources immediately available to it. This can be achieved by providing a standby coupling facility or spare capacity in other coupling facilities. The takeover procedure may require operator action, so requires operator training.

Coupling links are also a potential single point of failure. Multiple links should be provided for maximum availability.

In the event of a power failure that affects a coupling facility, a battery backup or uninteruptible power supply (UPS) can be provided to ensure that the contents of the coupling facility are not lost. This provides a faster recovery when power is restored.

In Figure 26 on page 88 we show a sample Parallel Sysplex configured for high availability. It has three processors and two coupling facilities. Each system is connected to both coupling facilities using two coupling links for redundancy. Both coupling facilities can be in use at once, but for different purposes. One could contain list and lock structures, while the other is used for cache structures.

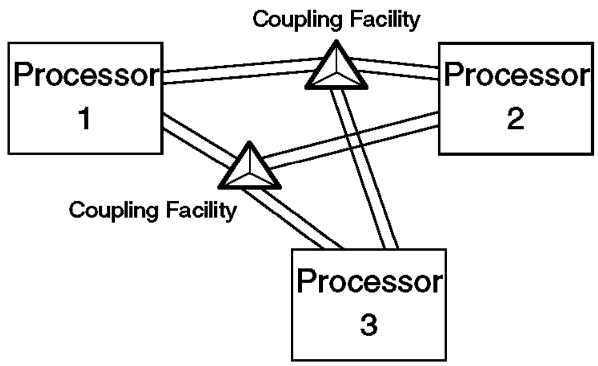

Figure 26. Backup of Coupling Facility

8.4 Sysplex Timer

The IBM 9037 Sysplex Timer synchronizes the Time-of-Day (TOD) clocks on up to 16 processors or processor sides. With the 9672, it is possible to connect more than one processor to each Sysplex Timer link; the R2 and R3 models can attach two processors to one Sysplex Timer link, and the E and P models can attach up to eight processors (CECs) to one link.

When running a sysplex with multiple processors, a common clocking mechanism is required. The 9037 allows events initiated by different processors to be time-stamped in the correct sequence.

Only one Sysplex Timer is required. However the Sysplex Timer is a possible single point of failure. If the Sysplex Timer fails, all MVS images in the sysplex will enter a non-restartable wait state, because proper event recording is no longer possible.

Once the sysplex is in production use, we therefore recommend that you implement two Sysplex Timers linked together, to provide redundancy. Each Sysplex Timer is attached to each processor side. One timer is active, the other is in stand-by mode. If the active Sysplex Timer or a processor attachment cable fails, time synchronization is maintained using the backup Sysplex Timer. No operator intervention is required. Figure 27 on page 89 shows the recommended configuration. The two Sysplex Timers must be within three meters of each other.

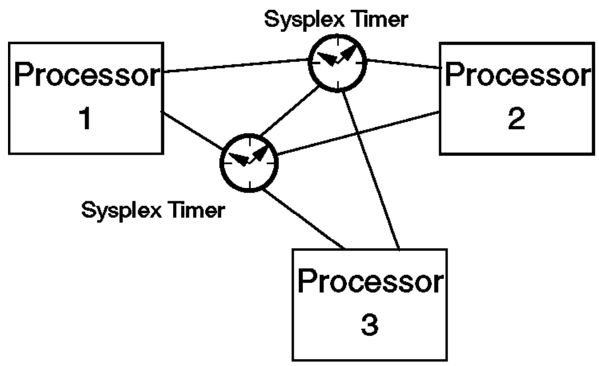

Figure 27. Backup of Sysplex Timer

8.5 I/O Connectivity

The implementation of a Parallel Sysplex may include the addition of more processors, and these will require access to shared DASD to be able to share data. It is therefore necessary to consider DASD connectivity, and indeed connectivity to all I/O devices. This section discusses hardware connectivity for I/O devices: DASD, tapes, network and consoles.

8.5.1 DASD

The only requirement for Parallel Sysplex is that the necessary connections between processors and DASD can be made. Neither ESCON nor IBM 3990 Model 6 controllers are required, though they will provide improved connectivity in larger Parallel Sysplex configurations.

If every processor is connected to every device, then more value is derived from the Parallel Sysplex, because all work can run everywhere. However, that may not be easy within your current configuration. In that case, consider having some processors that do not have access to all the DASD, by segmenting the workload and data. For example, one processor may run only development work and only needs access to development DASD, plus perhaps a small number of shared DASD for key data sets, such as the RACF security database. Even for transaction processing, the CICS file owning region (FOR) concept means that only one CICS region needs to access the databases for transactions running in other CICS regions. This does not provide all the benefits of Parallel Sysplex, but may be an effective interim solution until the DASD subsystem is upgraded. In Figure 28 on page 90 we show a simplified logical picture of this.

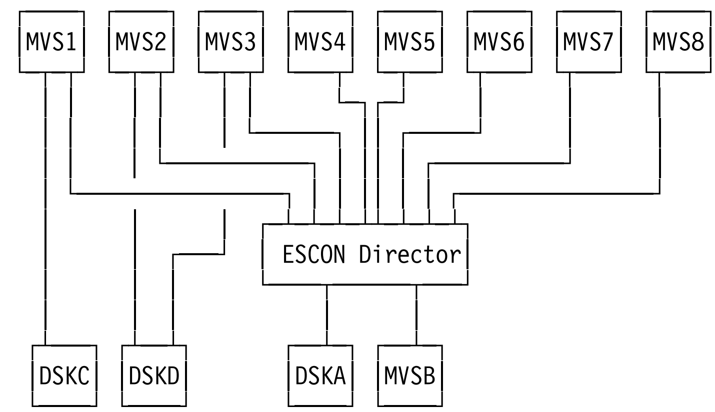

Figure 28. A Logical View of Disk Connections

Disks A and B would contain the data that must be shared, such as RACF and couple data sets. Disks C and D are only connected to a subset of the systems, and might contain page data sets, or specific application data.

8.5.1.1 ESCON

Parallel Sysplex does not require ESCON. In the Parallel Sysplex, the main requirement comes when you implement more than four MVS systems. If the Parallel Sysplex (together with other MVS systems not in the sysplex) has four or fewer MVS systems, then parallel channels can be used. With more than four systems, it becomes difficult to connect all the MVS systems into all the devices with full redundancy using parallel channels. One approach is to have fewer connections from some MVS systems. Four channel paths per MVS system is the usual number for a 3990 controller, but some systems may have sufficient capacity and redundancy with fewer than that number.

ESCON makes the connection of lots of MVS systems so much easier. To avoid problems with DASD connectivity in the future, we recommend that all future DASD be connected on ESCON channels, and on controllers that will provide enough logical channels for future requirements.

8.5.1.2 IBM 3990 Model 6

The IBM 3990 Model 6 DASD control unit is not required for Parallel Sysplex. However, even with ESCON, the IBM 3990 Model 3 will only support 16 logical channels. That means four channels to each of four MVS systems. If there are no more than four MVS systems, or full four channel connectivity from each MVS system is not necessary, then the 3990 Model 6 is not required. The 3990 Model 6 is needed to go beyond these limits. If you do implement the IBM 3990 Model 6, you will benefit from its superb performance, even if you don′t need its connectivity.

8.5.2 Tape

MVS/ESA SP Version 5.2 changes the way tape devices are selected for allocation. It considers the type of request, unit information on the request, and the characteristics of each available tape device. Based on these factors, it chooses the optimal device to allocate. Therefore, with MVS/ESA SP Version 5.2 an installation might notice a difference in which devices the system selects for a particular job.

Also new in MVS/ESA SP Version 5.2 is automatic tape switching, providing the ability for an installation to have a category of tape devices known as *automatically switchable*. These devices are not dedicated to any one system; rather they can be used by all systems in a Parallel Sysplex. Automatic tape switching improves the management of magnetic tape devices for systems in a Parallel Sysplex. It helps an installation to control costs by reducing both the number of tape devices and the amount of operator intervention required to manage the tape devices.

Prior to MVS/ESA SP Version 5.2, you need to consider which processors tape drives are connected to, and provide either manual switching or switching using a non-IBM software product. With MVS/ESA SP Version 5.2, you need to ensure that your tape devices support the dynamic switching capability.

8.5.3 Network

As you start your migration to Parallel Sysplex, it is important to note that your traditional subarea network can exist within the sysplex with little or no change. Your decision to implement APPN has no bearing on when you decide to implement your Parallel Sysplex. APPN is not required initially, and when introduced it can safely coexist with your traditional subarea network. APPN is required for the VTAM generic resource function. Generic resource provides single system image and session balancing for members of the generic resource. These functions are exploited by CICS.

As each new system is added to the Parallel Sysplex, access to the network through Front End Processors (FEP) is required. This can be achieved by ensuring your FEPs have sufficient channel adapters to support the increased connectivity requirements that will be required over time. Be aware of the limitations that your current hardware may place on you. If you are using 3705s or 3725s, note that the number of channel adapters can quickly become a constraint. 3745s can offer relief in this area when properly configured using an ESCON frame along with an ESCON Director. An alternative is to implement ESCON channel-to-channel connections (CTCs) to support VTAM to VTAM communication between MVS systems.

You may take this opportunity to review how you connect your network to the System/390 processors. As well as 3745/3746 connections, consider using 3172 or OSA adapters in your S/390 processors. You could connect a number of 9672s together with OSA adapters and a LAN. This topic is beyond the scope of this book. If you are interested in reviewing your network configuration then we recommend you discuss it with an IBM networking specialist.

8.5.4 Consoles

The implementation of a Parallel Sysplex is an opportunity to review your console configuration. Within a Parallel Sysplex there is only one sysplex master console (although multiple consoles may have master authority). This console is shared across all MVS systems, or at least it can be. When MVS systems are part of a Parallel Sysplex, there is a maximum of 99 MCS consoles, including both traditional MCS consoles and subsystem consoles such as those used by many subsystems. (Some subsystems use extended MCS consoles. The number of extended MCS consoles is not restricted).

Given this new limitation, does this mean that prior to migrating to a Parallel Sysplex you must change the way system consoles are physically cabled today?

No. However it does mean that as new MVS images are added to the complex, additional channel attached MCS consoles are not required. New MVS images can be set up to connect to the pool of consoles that already exist. This provides hardware savings as your Parallel Sysplex grows.

We recommend that you name your MCS and subsystem consoles. If you do not use console names, MCS assigns the next available console ID for the console whenever it rejoins the sysplex. For example, an unnamed console could be assigned a console ID of 05 in a sysplex with four consoles already active (IDs 01, 02, 03, and 04). If the system leaves the sysplex and rejoins later, MCS does not reassign the console ID 05 to the console, but instead assigns the next available ID, which is 06. Even if you have only five consoles in the sysplex, MCS assigns the next available ID. Sooner or later, MCS will have used all available 99 console IDs; then no new nameless consoles can be defined, and you have to bring the entire sysplex down to make new console IDs available.

With named consoles, MCS assigns the first available console ID to a console when it is defined to MCS for the first time. If the named console is redefined during a subsequent IPL, the same console ID is associated with the console, as long as the entire sysplex has not been IPLed.

Chapter 9. Creating Your Migration Plan

In the previous chapters we explained how Parallel Sysplex can help you meet your IT requirements. We looked at your applications to help you decide which ones should exploit Parallel Sysplex first. Then we looked at what that means to your operations, software and hardware.

The next step is to create a migration plan to implement your chosen application in a Parallel Sysplex. Your plan needs to fit in with your existing plans and commitments. This may be an iterative process. Having chosen an application, you may discover that it requires a set of software and hardware changes that are incompatible with your business, application, software migration, or capacity growth plans for the next 12 months. If so, go back to application selection and choose a more appropriate or simpler application to start with. This flow is illustrated in Figure 1 on page xiv.

Much of what you have to do to implement a Parallel Sysplex is “Business as Usual.” Installing new hardware and software, and updating systems management procedures, are tasks that you do regularly. You may already have plans to do these tasks in the next 12 months, in which case the additional effort of the migration to Parallel Sysplex may be minimal. There is a need to tie together Parallel Sysplex and existing software migration plans.

Some tasks in implementing a Parallel Sysplex are new, for example the implementation of the coupling facility and data sharing. These tasks are not difficult, and you can obtain assistance from IBM documentation and from IBM implementation services.

This chapter suggests a way that you can pull together the migration actions needed for your installation, and helps you to put these into a migration plan. In 9.1, “A Sample Migration Plan” on page 94 we provide a sample migration plan with many actions filled in. In A.8, “Blank Implementation Plan” on page 104 we provide you with a blank version of the plan. We suggest that you make a copy of it and use it to create your own implementation plan. To do that you need to do the following:

1. Mark on the plan the items that you already have in plan. You may have to adjust them slightly. Fill in when you were planning to do them.
2. Then fill in the dates for the items on the plan you were not already planning to do. Fit them in at appropriate times. This will enable you to decide what changes you need to make to your existing plan.
3. There are some blank lines in the plan for you to fill in additional items that you were either planning to do and that fit in with Parallel Sysplex implementation, or items you have identified from reading this book that you need to do and that are not on the sample plan.

9.1 A Sample Migration Plan

Sample Parallel Sysplex Migration Plan	1996			
	1	2	3	4
	Application (Chapter 4)			
Develop detailed plans				
Determine workload sizes and uniprocessor power requirements				
Run transaction affinity tool				
Adjust for affinities if required				
Split AORs/TORs if required				
Move a workload to new processor				
Establish a two system test environment for Parallel Sysplex				
Implement data sharing for first application				
Application and system test				
Implement transaction routing for first application				
Stress test first application				
Migrate first application to production				
	Operations (Chapter 6)			
Test coupling facility & develop procedures				
Update operations automation				
Train operators and update procedures				
	Software (Chapter 7)			
Systems programmer education				
Install and implement HCD				
Install MVS V5 & associated base software				
Install VTAM V4.2				
Install non-IBM software				
Install CICS V4.1				
Install IMS V5.1 and IRLM V2.1				
Implement test Parallel Sysplex with LPAR coupling facility				
Qualify for PSLC pricing				
Install CICSPlex/SM				
Network definition changes				
Implement APPN for generic resources				

Sample Parallel Sysplex Migration Plan	1996			
	1	2	3	4
Implement WLM goal mode				
Implement production Parallel Sysplex				
	Hardware (Chapter 8)			
Install ESCON Director and adapters				
Install processor upgrades or additional processors				
Install Sysplex Timer				
Install coupling links				
Reconfigure network as required				
Install production coupling facility				

9.2 Next Steps

Once you have a draft plan that illustrates what you could do in the next 12 months, your next steps should be:

- Validate that the plan will work technically. You need to validate that the application you have decided to move first into your Parallel Sysplex is appropriate. It should be appropriate in terms of its ability to data share, its ability to workload balance, and its capacity needs. IBM has tools and services to help you do that.
- Validate the plan financially. Depending on the choices you have for the hardware implementation, you need to review which are the financially attractive options. You also need to review your software and select a path that optimizes your expenditure. IBM is able to help you with this.
- Develop a detailed action plan. You should turn your draft plan into a detailed action plan. IBM has services to help you do that.

Appendix A. Worksheets

These worksheets are designed to be completed as you read through the book. The worksheets will help guide your thinking, and will enable you to record your thoughts for future reference.

We suggest that you make a copy of the worksheets, and complete the relevant worksheet after reading each chapter.

A.1 IT Requirements

This worksheet relates to the material in Chapter 1, “Identify Your Business Requirements” on page 1.

The rows list possible requirements that you may have. Fill in column three with your priority for each requirement. Use:

H for high priority
M for medium priority
L for low priority

We have also provided some blank rows for you to fill in any other requirements that are not listed here.

Table 4. Worksheet: IT Requirements, and How Parallel Sysplex Relates to Them

Requirement		Priority H/M/L	Data Sharing	Dynamic Workload Balancing	Single System Image	Cost Effective Hardware	Cost Effective Software
Application	Performance		√	√			
	Availability		√	√			
	Growth		√	√			
Overall	Performance		√	√			
	Availability		√	√			
	Growth		√	√	√		
Cost	Resource utilization		√	√			
	Hardware				√	√	
	Software						√
	People		√	√	√		
Additional items							

Note: The first three columns are explained in Chapter 1, “Identify Your Business Requirements” on page 1. The other columns are explained in Chapter 2, “How Parallel Sysplex Can Meet Your Requirements” on page 7. A

√ shows the primary aspects that will help you meet each requirement. Note that many of the requirements are interlinked; for example, if you make better use of your current hardware, you will delay your next upgrade and thus minimize your hardware and software costs. Therefore, use the table to show which aspects of Parallel Sysplex you wish to do *first*, while keeping in mind that you will probably wish to implement all aspects eventually.

A.2 Key Applications

This worksheet relates to the material in Chapter 1, "Identify Your Business Requirements" on page 1. In Table 5, list any key applications that you have identified, and check the requirements for each one.

Table 5. Worksheet: Key Applications

Application Name	Performance	Availability	Capacity Growth	Other (state)

A.3 Basic Migration Options

These questions relate to the material in Chapter 3, "Understanding Parallel Sysplex Migration Options" on page 15. Place a √ against the basic hardware configurations that may be relevant to you:

- Install a coupling-capable processor to replace an existing processor
- Install a Parallel Sysplex alongside an existing processor, not coupled to it
- Install a new CMOS processor and couple it to an existing processor
- Couple existing processors

A.4 Application Inventory

This worksheet relates to the material in Chapter 4, "Applications" on page 19 and Chapter 5, "Application Subsystem Support of Parallel Sysplex" on page 27.

Table 6. Worksheet: Application Description and Selection for Exploiting Parallel Sysplex

Application Suite Name (4.2.1)	**Business Benefit of Parallel Sysplex (4.2.2)**	**Risks (4.2.3)**	**Need ICRF or Vector Facility (4.2.4)**	**Total Capacity (4.2.5)**	**Uniprocessor Power (4.2.6)**	**Subsystem or Transaction Manager (4.2.7)**	**Database Manager (4.2.8)**	**Exploitation of Data Sharing (5.1)**	**Exploitation of Dynamic Workload Balancing (5.2)**	**Priority - H/M/L/X (5.3)**

A.5 Systems Management and Operational Considerations

These worksheets relate to the material in Chapter 6, "Systems Management and Operational Considerations" on page 45. Note here any changes you need to make, or wish to make.

A.5.1 Systems Management and Operation of a Multisystem Environment

A.5.2 Backup and Recovery in a Parallel Sysplex

A.6 Software Worksheets

These worksheets relate to the material in Chapter 7, "Software" on page 55. Note here any changes you need to make, or wish to make.

A.6.1 Software Packaging

A.6.2 Non-IBM Software

A.6.3 Minimizing the Migration Time and Effort

A.6.4 Testing the Changes

A.6.5 Software Maintenance Strategy

A.6.6 Changes in Software Pricing

A.6.7 Base Software

A.6.8 Database Manager Software

A.6.9 Transaction Manager Software

A.6.10 Networking and Systems Management Software

A.7 Hardware Worksheets

These worksheets relate to the material in Chapter 8, "Hardware" on page 77. Note here any changes you need to make, or wish to make.

A.7.1 Draw a Picture of Your Current Processors

A.7.2 Draw Pictures of Possible Hardware Changes

Include the following items as appropriate:

- What coupling facility options are relevant, for test and production?
- What coupling facility links (number and type) do you need?
- How many Sysplex Timers are required?
- How much ESCON is required and when?
- What changes must take place within my network to support the Parallel Sysplex?
- How should I configure my consoles?

A.8 Blank Implementation Plan

This plan relates to the material in Chapter 9, “Creating Your Migration Plan” on page 93. It is a blank copy of the sample migration plan in that chapter. Blank lines have been left in the plan for you to add additional tasks for your implementation.

Your Parallel Sysplex Migration Plan	**1996**			
	1	**2**	**3**	**4**
	Application (Chapter 4)			
Develop detailed plans				
Determine workload sizes and uniprocessor power requirements				
Run transaction affinity tool				
Adjust for affinities if required				
Split AORs/TORs if required				
Move a workload to new processor				
Establish a two system test environment for Parallel Sysplex				
Implement data sharing for first application				
Application and system test				
Implement transaction routing for first application				
Stress test first application				
Migrate first application to production				
	Operations (Chapter 6)			
Test coupling facility & develop procedures				
Update operations automation				
Train operators and update procedures				

Your Parallel Sysplex Migration Plan	1996			
	1	2	3	4
	Software (Chapter 7)			
Systems programmer education				
Install and implement HCD				
Install MVS V5 & associated base software				
Install VTAM V4.2				
Install non-IBM software				
Install CICS V4.1				
Install IMS V5.1 and IRLM V2.1				
Implement test Parallel Sysplex with LPAR coupling facility				
Qualify for PSLC pricing				
Install CICSPlex/SM				
Network definition changes				
Implement APPN for generic resources				
Implement WLM goal mode				
Implement production Parallel Sysplex				
	Hardware (Chapter 8)			
Install ESCON Director and adapters				
Install processor upgrades or additional processors				
Install Sysplex Timer				
Install coupling links				
Reconfigure network as required				
Install production coupling facility				

Appendix B. Additional Material for IBM Employees

This chapter supplements "Related Publications" on page xv with a list of additional publications, documents and tools available to IBM employees. These are suitable for a more detailed discussion.

B.1.1 International Technical Support Organization Publications

- *ESA/390 Coupling Technologies Presentation Guide*, ZZ81-0340
- *Parallel Sysplex Presentation Guide Systems Management*, ZZ81-0341
- *MVS 5.1 Workload Manager Presentation Guide*, ZZ81-0335

B.1.2 IBM Tools Disks

The following packages are all available on MKTTOOLS. They can be accessed locally if available with the command "TOOLCAT MKTTOOLS," or they can be accessed from the central repository by typing the following command from your VM user ID:

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET pname PACKAGE

where "pname" is the name of the package you want.
```

In this list we give the package name in uppercase letters, followed by the package title:

PSNOW *System/390 MVS Parallel Sysplex - Start the Migration Now*, the ITSO redbook ZZ81-0354

PSCOMPAR *A Comparison of System/390 Configurations - Parallel and Traditional*, the ITSO redbook SG24-4514 and a foil presentation

PTSVALUE *Customer Value While Migrating to Parallel Sysplex*

PTSMIGR *Parallel Sysplex Migration: Start to Finish* foil presentation

PTSMIG *PTS Migration Guide*

PTSCVW *S/390 Parallel Sysplex Offering Customer Value Workshop Guide*

MFTSPERF *Performance for S/390 Parallel Transaction Server*

LSPRPC *Large Systems Performance Reference (LSPR)* PC tool

SC281187 *Large Systems Performance Reference (LSPR)* publication

BWATOOL *Batch Workload Analysis Tool*

BWA2OPC *Predicting Batch Turnaround when Moving to Smaller Engines*

CICSPTS *CICS in a Parallel World*

CICSE41P *CICS/ESA 4.1 Performance Comparison With CICS/ESA 3.3*

CICSCAU1 *CICS Transaction Affinities Utility MVS/ESA Presentation*

CICSVSAM *CICS and VSAM: Data Sharing in an MVS Environment*

PTSHELP *A Directory of Parallel Resources*, including a list of Parallel Sysplex books and deliverables

Glossary

ACF/VTAM. Advanced Communications Function for the Virtual Telecommunications Access Method. See *VTAM*.

affinity. A connection or association between two objects.

application. A collection of software components used to perform specific types of work on a computer, for example, a payroll application, an airline reservation application.

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.

cryptographic. Pertaining to the transformation of data to conceal its meaning.

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.

dynamic. Pertaining to an operation that occurs at the time it is needed rather than at a predetermined or fixed time.

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

ICMF. Integrated Coupling Migration Facility. It runs on IBM 9021 711-based processors, 9121 511-based processors, and the 9672. It provides coupling facility functions for MVS/ESA images running in logical partitions on the same processor complex.

ICRF. Integrated Cryptographic Feature. An optional facility on certain IBM processor complexes for encrypting data more rapidly. See *cryptographic*.

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

IPL. Initial Program Load. The initialization procedure that causes an operating system to start operation.

ITSO. International Technical Support Organization

JCL. Job control language.

LAN. Local area network. A data network located on the user′s premises in which serial transmission is used for direct data communication among data stations. It services a facility without the use of common carrier facilities.

logical partition. In LPAR mode, a subset of the processor unit resources that is defined to support the operation of a system control program (SCP).

LPAR. Logically partitioned mode. A mode that allows the operator to allocate hardware resources of the processor unit among several logical partitions.

MVS. Multiple Virtual Storage operating system.

Parallel Sysplex. A sysplex that uses one or more coupling facilities.

processor. A physical collection of hardware that includes main storage, one or more processor engines, and channels.

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

PTF. Program temporary fix.

RACF. Resource Access Control Facility. An IBM licensed program that provides for access control by identifying and verifying users to the system, authorizing access to DASD data sets, logging detected unauthorized attempts to enter the system, and logging detected accesses to protected data sets.

RMF. Resource management facility

SEC. System engineering change.

shared data. See *data sharing*.

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.

TCP/IP. Transmission control protocol/internet protocol. A public domain networking protocol with standards maintained by US Department of Defense to allow unlike vendor systems to communicate.

TSO. Time sharing option.

uniprocessor. One of the processor engines in a processor.

vector facility. An optional facility on certain IBM processor complexes for processing vector calculations more rapidly.

VSAM. Virtual Storage Access Method. An access method for direct or sequential processing of fixed and variable-length records on direct access devices.

VTAM. Virtual telecommunication access method. This program provides for workstation and network control. It is the basis of a System Network Architecture (SNA) network. It supports SNA and certain non-SNA terminals. VTAM supports the concurrent execution of multiple telecommunications applications and controls communication among devices in both single-processors and multiple processors networks.

XCF. cross-system coupling facility. XCF is a component of MVS that provides functions to support cooperation between authorized programs running within a sysplex.

Index

Numerics

3172 91
3705 91
3725 91
3745 91
3990 89, 90
9037 Sysplex Timer 88
9672 5, 7, 13, 18
9674 Coupling Facility 82

A

abbreviations 109
ACF/TCAM 74
acronyms 109
ADAPLEX+ 31, 59, 71
affinity 36—38, 39, 56
AOC/MVS 43, 76
application
- affinity 36, 56
- availability 3, 22
- business value 2, 21
- capacity 24
- CICS macro level 36, 72
- clone 73, 74
- current 18
- inventory 20, 21, 99
- map to systems 42
- migration 19
- new 23
- performance 3
- priority 41
- selection 20
- size 3
- split 10

APPN 75, 91
Automatic Restart Manager (ARM) 50, 68, 69
automation, console 47
availability
- application 3, 8, 22, 48
- coupling facility 83, 87
- coupling link 87
- DASD 58
- data sharing 10
- dynamic workload balancing 11
- Sysplex Timer 88
- system 4

B

backup
- coupling facility 51, 87
- coupling link 87
- disaster 51, 56

backup *(continued)*
- procedures 56
- processor 49, 50
- Sysplex Timer 51, 88

batch 25, 26, 40, 74
batch workload analysis tool 107
Batchpipes 26
BMP 41, 74
business benefit 21
business requirements 1
business risk 23
BWATOOL 107

C

capacity planning 24, 25, 50, 81
capacity, uniprocessor power 25, 81
catalog, master 58
CEC level pricing 60
CFCC 82, 83
CICS
- affinity 36, 37
- dynamic workload balancing 32, 34
- Intersystem Communication (ISC) 33, 35
- macro level 36, 72
- Measured Usage License Charge 61
- multiregion operation (MRO) 33, 35, 37, 68, 72
- Parallel Sysplex exploitation 68, 72
- Parallel Sysplex Licence Charge 59
- performance 38, 72
- systems management 73
- temporary storage 31
- terminal owning region (TOR) 37
- Transaction Affinities Utility 36, 107
- transaction routing 33
- uniprocessor power 34
- Version 2 72
- VSAM record level sharing 31, 65
- with IMS/DB 29
- workload distribution 32

CICSPlex/System Manager 32, 35, 73
clone application 73, 74
CMC 75
CMOS 5, 7, 13, 80
connectivity, DASD 89, 90
console 46, 47, 67, 68, 91
continuous operation 22
cost of computing
- granularity 4
- hardware 5, 8, 13, 52
- lease charges 4
- people 6, 11, 12, 13, 48
- resource utilization 5, 8, 12
- software 5, 8, 13, 58, 83
- total 4

coupling facility 51, 82, 86, 87
coupling facility control code (CFCC) 82, 83
coupling link 86, 87
cryptographic feature 24, 36, 81

D

DAE 67, 68
DASD 89
DASD, RAMAC 58
DASD, shared 9, 46, 89
data sharing 8—11, 27—31, 69
database manager 69—71
DB2 30, 59, 61, 71
DBCTL 28, 70, 72
DBRC 28, 70
DFSMS/MVS 59, 60, 62, 65, 67, 68
direct access method (DAM) 31
disaster recovery 51, 56
DL/I 28
dynamic workload balancing 8, 11—12, 31—41, 44

E

education, operator 49
ESCON 89, 90
ESCON Manager 65

F

failure recovery 49
fiber, coupling link 86

G

generic resource 35, 69, 75, 91
glossary 109
goal mode 66
Graduated License Charge 59
granularity 10, 22
growth 3, 4, 8, 22
growth, data sharing 10
growth, dynamic workload balancing 12
growth, single system image 13
GRS 9, 65, 67, 68, 69

H

hardware
  console 91
  cost 5
  coupling facility 82
  DASD 89
  for Parallel Sysplex 77
  I/O connectivity 89
  network connection 91
  Parallel Sysplex Licence Charge (PSLC) 77
  processor 78
  Sysplex Timer 88
hardware *(continued)*
  tape 90
Hierarchical File System (HFS) 31

I

I/O connectivity 89
IBM services 95
ICMF 85
ICRF 24, 36, 81
IMLC 61
IMS BMP 41, 74
IMS multiple system coupling (MSC) 39, 73
IMS/ESA database manager 28, 59, 61, 70
IMS/ESA transaction manager 28, 38, 59, 61, 68, 73
IMS/TM affinity 39
Indexed Monthly License Charge (IMLC) 61
Info/Management 43
Integrated Coupling Migration Facility (ICMF) 85
Integrated Cryptographic Feature (ICRF) 24, 36, 81
inventory, application 99
IRLM 28, 68, 70
ISC 33, 35
isolate workload 42
IT requirements 97

J

JES2 64, 68
JES3 64, 68, 69

L

LAN management/serving 43
Large Systems Performance Reference (LSPR) 24, 107
link, coupling 86, 87
local DL/1 28, 72
logging 68
LPAR 83
LSPR 24, 107

M

macro level, CICS 36, 72
maintenance, software 57
maintenance, system 48
master catalog 58
MCS console support 69
Measured Usage License Charge (MULC) 61
migration
  effort 56
  gradual 19
  options 15, 98
  plan 93, 94, 104
  software 55
mono-mode fiber link 87
MRO 33, 35, 37, 68, 72

MRO performance 38, 72
MSC 38, 39, 73
MULC 61
multi-mode fiber link 86
multiple systems coupling (MSC) 38, 39, 73
multiregion operation (CICS MRO) 33, 35, 37, 68, 72
MVS (base) software 62
MVS JES2 64, 68
MVS JES3 64, 68, 69
MVS workload manager 66, 68

N

naming conventions 57
NetView 46, 47, 74, 76
network 91
network management 43
networking software 74
non-IBM software 55

O

OPC/ESA 43, 76
OpenEdition Hierarchical File System (HFS) 31
operations 45, 46
operations, console 47
operator training 49
OSA adapter 91

P

Parallel Sysplex
  description 7
  hardware 77
  migration
    See migration
  operations 45
  performance 107
  software required 62
  systems management 45
Parallel Sysplex Licence Charge (PSLC) 5, 13, 58, 59, 77
parameters, symbolic 57
partitioned access method (PAM) 31
people cost 6, 11, 12, 13, 48
performance
  application 3, 8
  CICS 38, 72
  data sharing 9
  dynamic workload balancing 11
  management 47
  Parallel Sysplex 107
  system 3
  testing 56
peripherals 45
plan, migration 93, 94, 104
printing 43
priority 41
processor
  backup 50
  recovery 49
  selection 78
PSLC 5, 13, 58, 59, 77

R

RACF 48, 68, 69, 76
RAMAC DASD 58
record level sharing (RLS) 31, 65
recovery
  See backup
requirements 1, 8, 20, 97
reserve/release 9
resource utilization 5, 8, 10, 12
response time 25
risk 23
RMF 47, 68, 69

S

security management 43
sequential access method (SAM) 31
services, IBM 95
shared DASD 9, 46, 89
shared SYSRES 48
shared tape 66, 69, 91
sharing peripherals 45
single system image 6, 8, 12—13, 46
SLIP 67, 68
SMS 65
software
  cost 5, 58, 83
  maintenance 57
  migration 55
  networking 74
  non-IBM 55
  packaging 55
  systems management 76
Software AG 31, 71
switching peripherals 45
symbolic parameters 57
Sysplex Timer 51, 88
SYSRES 48, 57, 58
system maintenance 48
system managed storage 65
systems management 6, 11, 12, 45
systems management software 43, 76
systems management, CICS 73

T

tape sharing 66, 69, 91
TCP/IP 74, 75
temporary storage, CICS 31
terminal owning region (TOR) 37
testing 56

training, operator 49
transaction affinity 36
transaction manager 71—74
transaction routing, CICS 33
TSCF 76
TSO 26, 40, 61, 67, 68, 69, 74

U

uniprocessor power 25, 34, 81

V

vector facility 24, 36, 81
VLF 67, 68
VSAM record level sharing 31, 65
VTAM
  APPN 75
  CMC 75
  definitions 75
  generic resource 35, 69, 75, 91
  Parallel Sysplex exploitation 68, 69
  Parallel Sysplex Licence Charge 59, 60
  supported levels 74

W

workload balancing 11, 44
workload distribution 32, 38, 42, 44, 48
workload isolation 42
workload manager 66, 68
workload selection 20
worksheet 97

X

XCF 67, 69
XES 67

ITSO Technical Bulletin Evaluation RED000

International Technical Support Organization
System/390 MVS Parallel Sysplex
Migration Paths
February 1996

Publication No. SG24-2502-00

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**
SG24-2502-00

IBM®

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.

SG24-2502-00