


Independent ASP Performance Study on the IBM *e*server iSeries Server

Understand how IASP drive configuration
can affect your application

Understand user profile mismatch
implications

Learn about an SAP true life
implementation



Hernando Bedoya
Emilio Cortes
Manfred Engelbart
Vivian Kirkpatrick
Michael Lance



International Technical Support Organization

**Independent IASP Performance Study on the
IBM @server iSeries Server**

April 2004

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (April 2004)

This edition applies to Version 5, Release 2, Modification 0 of OS/400.

© Copyright International Business Machines Corporation 2004. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	v
Trademarks	vi
Preface	vii
The team that wrote this Redpaper	vii
Become a published author	viii
Comments welcome	ix
Part 1. Background	1
Chapter 1. Introduction to Independent ASP	1
1.1 Introduction/overview	2
1.2 Definition of terms or terminology	3
1.3 IASP possible configurations	4
1.3.1 Single system environment	4
1.3.2 Multisystem clustered environment	7
1.4 Performance requirements	8
Chapter 2. Testing environments	9
2.1 Cluster hardware configurations	10
2.1.1 Customer SAP environment	10
2.1.2 Lab testing environments	10
2.2 Application software	16
2.2.1 SAP - Customer scenario	16
2.2.2 SQL datawarehousing queries in a banking environment	16
2.2.3 SQL-based benchmark testing application	17
2.2.4 RPG-based benchmark testing application	17
2.3 OS/400	18
2.3.1 User profiles	18
2.3.2 Failover times	21
2.3.3 Switchover times	22
Part 2. Performance findings	25
Chapter 3. SAP - Customer scenario	27
3.1 Introduction - The company profile VEKA AG	28
3.2 VEKA's solution overview	28
3.2.1 SAP architecture overview	28
3.2.2 The original installation	28
3.2.3 The new cluster-based solution design	31
3.3 SAP installation overview	33
3.4 The cluster environment	35
3.4.1 The device CRG	38
3.4.2 The application CRG	39
3.5 Conclusions and findings	41
3.5.1 Failover/switchover times	41
3.5.2 I/O performance considerations	42
3.5.3 Disk layout considerations	44
3.5.4 Transaction profile and total response time	47

Chapter 4. SYSBAS vs IASP configurations	49
4.1 Why SYSBAS vs IASP configurations are important	50
4.2 Different DASD configuration scenarios	50
4.2.1 ITSO lab environment - iSeries Model 810	50
4.2.2 Cluster lab environment - iSeries Model 825	50
4.3 Results	51
4.3.1 ITSO Lab - iSeries Model 810 in a testing application	51
4.3.2 Cluster Lab - iSeries Model 825 in a testing application	55
4.3.3 Another view with the total number of disk arms constant	60
4.3.4 ITSO Lab - iSeries Model 810 in a data warehousing banking environment	66
4.3.5 Cluster Lab - iSeries Model 825 in a Data warehousing Banking Environment	71
4.4 Observations and comments	75
Chapter 5. User profile considerations	77
5.1 Why are user profiles important to IASP Performance	78
5.2 Data used for the User Profile testing	79
5.3 Failover vs. switchover	82
5.4 Comparing UID vs. UID/GID profile mismatches	88
5.5 Library vs. IFS	89
5.6 How much difference a processor makes	92
5.7 OK, how do I fix it	93
Part 3. Conclusions and recommendations	97
Chapter 6. Conclusions and recommendations	99
6.1 General performance findings	100
6.2 Other performance recommendations	101
Appendix A. Supported and unsupported objects in IASPs	103
Supported object types	104
Unsupported object types	104
Related publications	105
IBM Redbooks	105
Other publications	105
Online resources	105
How to get IBM Redbooks	105
Help from IBM	105

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law. INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

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

AS/400®	Enterprise Storage Server®	Redbooks™
CICS/400®	IBM®	Redbooks(logo)  ™
Distributed Relational Database Architecture™	ibm.com®	System/38™
Domino®	iSeries™	xSeries®
DB2®	Lotus®	
	OS/400®	

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Preface

Independent ASPs were announced and made available in V5R1. The initial introduction of IASPs was limited to IFS objects. As such, iSeries™ customers started using this concept for Domino®, Web serving applications, and Integrated xSeries® Adapters servers. In V5R2 some major enhancements were made to the support of IASP. Now they support traditional library and object structure in the QSYS.LIB file system used by the majority of iSeries applications. For this reason many customers have been asking what the performance considerations are when you move your application to IASPs.

The aim of this Redpaper is to provide you with an understanding of the performance implications when you decide to move your applications to Independent ASP. This redpaper should be read in conjunction with the existing IBM® Redbook *IBM eServer iSeries Independent ASPs A Guide to Moving Applications to IASPs*, SG24-6802. It is the purpose of this redpaper to complement the existing IBM Redbook using real customer scenarios to illustrate some performance considerations.

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Rochester Center.

Hernando Bedoya is an IT Specialist at the IBM ITSO, in Rochester, Minnesota. He writes extensively and teaches IBM classes worldwide in all areas of DB2® UDB for iSeries. Before joining the ITSO more than one year ago, he worked for IBM Colombia as an AS/400® IT Specialist doing presales support for the Andean countries. He has 20 years experience in the computing field and has taught database classes in Colombian universities. He holds a Masters Degree in computer science from EAFIT, Colombia. His areas of expertise are database technology, application development, and data warehousing

Emilio Cortes is a computer engineer who is responsible of leading all the iSeries services and technology from Corpbanca(Bank) in Chile. He has ten years of experience working with AS/400 and iSeries systems. He has participated as system engineer, database administrator and system analyst. He got his first engineering degree in the Universidad de Concepción of Chile in 1986 and at the moment he is completing his fourth year to get his second engineering degree, this time in industries engineering, at the Universidad de Santiago de Chile.

Manfred Engelbart is an IT Specialist at the IBM Technical Sales Support team in the Central Region doing technical presales support for the SAP on iSeries solution. He has worked at IBM for 29 years, has 15 years experience with IBM iSeries, and has coauthored different redbooks and redpapers about SAP and i2 implementation on IBM iSeries.

Vivian Kirkpatrick is a Staff Software Engineer who works in the Rochester Support Center in the United States. He has nine years of experience on the iSeries and AS/400. He holds a degree in Computer Information Science from Minnesota State University - Mankato. He worked three years in Database support before moving to the SLIC/VMC/Internals area, where he has worked for three years. He has taught multiple courses for Rochester and World Wide support on Cluster and IASP.

Michael Lance is a Systems Engineer III who works at Countrywide Financial Corporation in Simi Valley, CA. He has 23 years of experience on System/38™, AS/400 and iSeries servers.

He is an IBM Certified Systems Expert - iSeries LPAR Technical Solutions V5R2. He attended Guildford Technical College in Guildford, England. He has written various copyrighted utilities that enhance OS/400® system management functionality with his company Lance Consulting.

Thanks to the following people for their contributions to this project:

Special thanks to the Cluster Lab Team in IBM Rochester:

Ron Peterson
Tim Covalt

Special thanks to the following customers for authorizing us to use their environment for testing purposes:

Countrywide Financial Corporation
Christopher Jones from Countrywide Financial Corporation
Corpbanca in Chile
VEKA AG from Germany

Michael Brinker
Gordie Grout
Gaylene Kirkpatrick
Mike Konke
Nick Lawrence
Timothy Mossing
Vicki Morey
Glen Nelson
Amartey Pearson
Travis Pizel
James Ranweiler
Steven Royer
Mike Snyder
Kiswanto Thayib
Craig Wilcox
Jon Wold
Judith Zavala
IBM Development

Christine Felner-Goesch
Michael Cain
Chuck Stupca
Selwyn Dickey
Ted Bauer
IBM Rochester

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this Redpaper or other Redbooks™ in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an Internet note to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. JLU Building 107-2
3605 Highway 52N
Rochester, Minnesota 55901-7829



Part 1

Background

In this part we:

- ▶ Discuss what IASPs are and why they are important to customers.
- ▶ Discuss why we wrote this paper.
- ▶ Describe some of the new V5R2 IASP enhancements.
- ▶ Describe the testing environment, hardware and software that was used for this project.



Introduction to Independent ASP

An Independent auxiliary storage pool (IASP), or independent ASP, is a collection of disk units that can be brought online or taken offline independent of the rest of the storage on a system, including the system ASP, user ASPs, and other IASPs. An IASP can be either:

- ▶ *Private*: Privately connected to a single system or system partition (in an Logically Partitioned environment), also known as stand-alone IASPs
- ▶ *Switchable*: Switched between two (or more with LPAR) systems or partitions in a clustered environment

Clearly, this is quite a departure from the way in which auxiliary storage (disk) was regarded prior to V5R1. Until then, all iSeries disks were considered to be owned and usable only by a single system. Enhancements made to OS/400 in V5R1 and again in V5R2 make using IASPs an attractive option for many customers who are looking for higher levels of availability and server consolidation.

This chapter gives an introduction to the terminology and concepts related to IASPs. It also describes the questions that we want to answer regarding IASP performance.

1.1 Introduction/overview

The concept of “independent storage” is new to the iSeries server. For many years, iSeries customers, AS/400, and System/38 customers before that, used a single view of storage, which was totally owned by the system to which it was attached.

User ASPs were introduced on System/38 to allow the disk storage attached to a single system to be grouped into separate pools. However, these pools always had a close relationship to the system ASP. You could not have duplicate library names in different user ASPs. In effect, all the storage attached to a system was a single name space. In V5R2 this type of User ASP is now called a Basic ASP. The term SYSBAS was created to describe non-Independent ASPs which are made up of the system ASP and all Basic ASPs as shown in Figure 1-1.

In more recent versions of OS/400, various file systems, such as QOpenSys, QNetWare, and QNTC, were introduced. They were generically known as the *integrated file system* (IFS). By defining user-defined file systems, the IFS objects could be stored in different user ASPs. However, these were limited to only available on the system to which the disk storage was attached.

In V5R1 of OS/400, the first move away from this single view of storage was introduced with the ability to create an Independent Auxiliary Storage Pool (IASP). An IASP is a collection of disk units that can be brought online or taken offline independent of the rest of the storage on a system. An IASP can be switched between systems or logical partitions. This initial introduction of IASPs was limited to storing only IFS objects. As such, iSeries customers had limited use of IASPs. IASPs could be used for Domino, Web serving applications, and Integrated xSeries Adapters or Servers. This ASP type is called an UDFS IASP.

With the enhancements introduced in V5R2 in the form of Primary and Secondary IASPs, we have support for the traditional library and object structure in the QSYS.LIB file system used by the majority of iSeries applications. See Appendix A, “Supported and unsupported objects in IASPs” on page 103, for a full list of objects that are supported in an IASP and a list of objects that are not supported in an IASP.

IASPs introduce many new potential advantages to most customers. Among the benefits are:

- ▶ Availability
- ▶ Server consolidation of “branch office” type systems
- ▶ Workload balancing across multiple servers through switching
- ▶ Reclaim storage by IASP
- ▶ Isolation of data associated with specific applications
- ▶ Save and restore management by IASP
- ▶ Perform application maintenance that does not affect the entire system

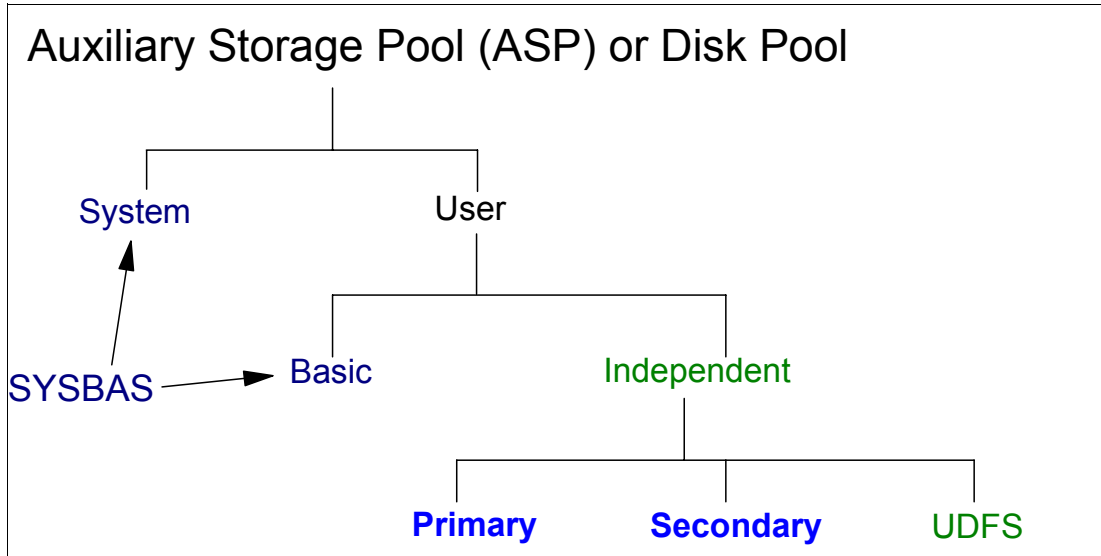


Figure 1-1 Auxiliary storage pool types

1.2 Definition of terms or terminology

This section briefly defines the following terms:

- ▶ *Auxiliary Storage Pool*: (aka, ASP) An auxiliary storage pool that contains only disk units.
- ▶ *Auxiliary Storage Pool Group*: (aka, Group or ASP Group) Made of a primary IASP and zero or more secondary IASPs, each of which are independent in regard to data storage, but combine to act as one entity.
- ▶ *Cluster*: A collection of complete systems that work together to provide a single, unified computing capability.
- ▶ *Cluster Resource Group (CRG)*: Cluster Resource Group (CRG): A collection of related cluster resources that defines actions to be taken during a switchover or failover operation of the access point of resilient resources. The group describes a recovery domain and supplies the name of the cluster resource group exit program that manages the movement of an access point. A device CRG contains a list of switchable devices, such as IASPs which reside on a switchable entity. A switchable entity can be either an expansion unit (tower) or an IOP. In iSeries Navigator, a device cluster resource group is referred to as a switchable hardware group.
- ▶ *Failover*: A cluster event where cluster resource group(s) switch over to a backup system due to the failure of the primary system.
- ▶ *HSL (high speed link)*: The system-to-tower connectivity technology that is required to implement switchable IASPs residing on an expansion unit (tower). The servers and towers in a cluster using resilient devices on an external tower must be on and HSL loop connecting with HSL cables.
- ▶ *Independent Auxiliary Storage Pool*: (aka, Independent ASP or IASP) One or more storage units that are defined from the disk units or disk-unit subsystems that make up addressable disk storage. An IASP contains objects, the directories and libraries that contain the objects, and other object attributes such as authorization ownership attributes. An IASP can be made available (varied on) and made unavailable (varied off) without restarting the system. An IASP can be either a) switchable among multiple systems in a clustering environment or b) privately connected to a single system.

- ▶ *Library name space*: An attribute that can be set for the current thread. The library name space is the set of objects and libraries that can be accessed in any IASPs in a ASP group plus the libraries in the system ASP and basic user ASPs (ASPs 2-32) using the regular library-qualified object name syntax. The Set Auxiliary Storage Pool Group (SETASPGRP) command sets the auxiliary storage pool (ASP) group for the current thread.
- ▶ *Primary Independent Auxiliary Storage Pool*: (aka, Primary IASP or Primary Disk Pool) An IASP that defines a collection of directories and libraries and may have other secondary ASPs associated with it. A primary ASP also defines a database for itself and other ASPs that may be added to its ASP group. Primary ASPs can only be implemented on V5R2 or later of OS/400.
- ▶ *Switchable entity*: The physical resource containing the IASPs that can be switched between systems in a cluster. This can be a expansion unit containing disk units in a multiple system environment. This could also be an IOP containing disk units in an LPAR environment.
- ▶ *Switchover*: A cluster event where the cluster resource group(s) switches over to a backup system due to the manual intervention from the cluster management interface.
- ▶ *SYSBAS*: This refers to the system ASP (ASP 1) and all configured basic ASPs (ASPs 2-32). IASPs (APSs 33-255) are not included.

1.3 IASP possible configurations

There are several possible configurations, but only two basic environments in which IASPs can be used:

- ▶ A single system environment with a single iSeries server or partition (LPAR)
- ▶ A multisystem (or multipartition) environment managed by an iSeries cluster

1.3.1 Single system environment

In a single system environment, you can take IASPs offline or make them unavailable, independent of other ASPs. You can also bring the IASP online or make it available, while the system is active, without performing an initial program load (IPL) of the whole system.

Figure 1-2 shows an example of multiple databases that reside in IASPs. This example has IASPs for the Payroll data, Order entry data, and data for Companies 1, 2, and 3. The actual application code could reside in the system ASP or another user ASP.

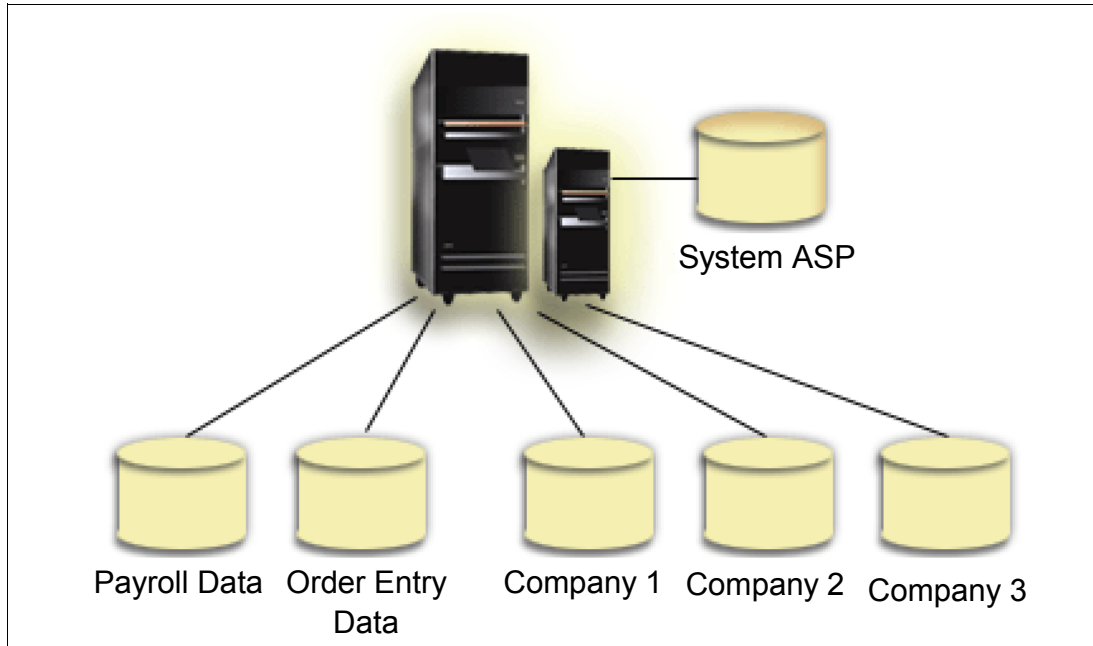


Figure 1-2 Single system IASP environment

A typical use of IASPs as shown in this example is for server consolidation of multiple branch offices or store systems. Corporate data can reside in the other IASPs. Segmenting your databases in this manner allows for greater control and flexibility.

Part of the attraction of using IASPs for multiple databases, is that you can have the same library names in multiple IASPs if they are part of different ASP groups. This is because of a new V5R2 function called namespaces.

Figure 1-3 shows a simple example of namespaces. The gentleman in the upper right has no ASP group in his namespace, so he can not use the libraries in the IASPs. The lady in the lower left corner has a namespace that allows her to access the system ASP as well as the ASP group Finance, so she can get into libraries in either place. The gentleman in the lower right corner can access libraries in the system ASP or the Sales ASP group.

Because of the namespace concept we can have libraries with the same name in the Sales ASP group as well as the Finance ASP group and the system will know which library the user would need to get to based on their namespace. For example if we have the library TEST in the Finance ASP group as well as the Sales ASP group; and all three users tried to get to the library SALES, here is what would happen:

- ▶ The upper right gentlemen would get a message telling him there is no library found. This is because he has no ASP group in his namespace, and the base system does not have a library named TEST.
- ▶ The lady in the lower left would get into the library TEST in the Finance ASP group.
- ▶ The gentleman in the lower right would get into the library TEST in the Sales ASP group.

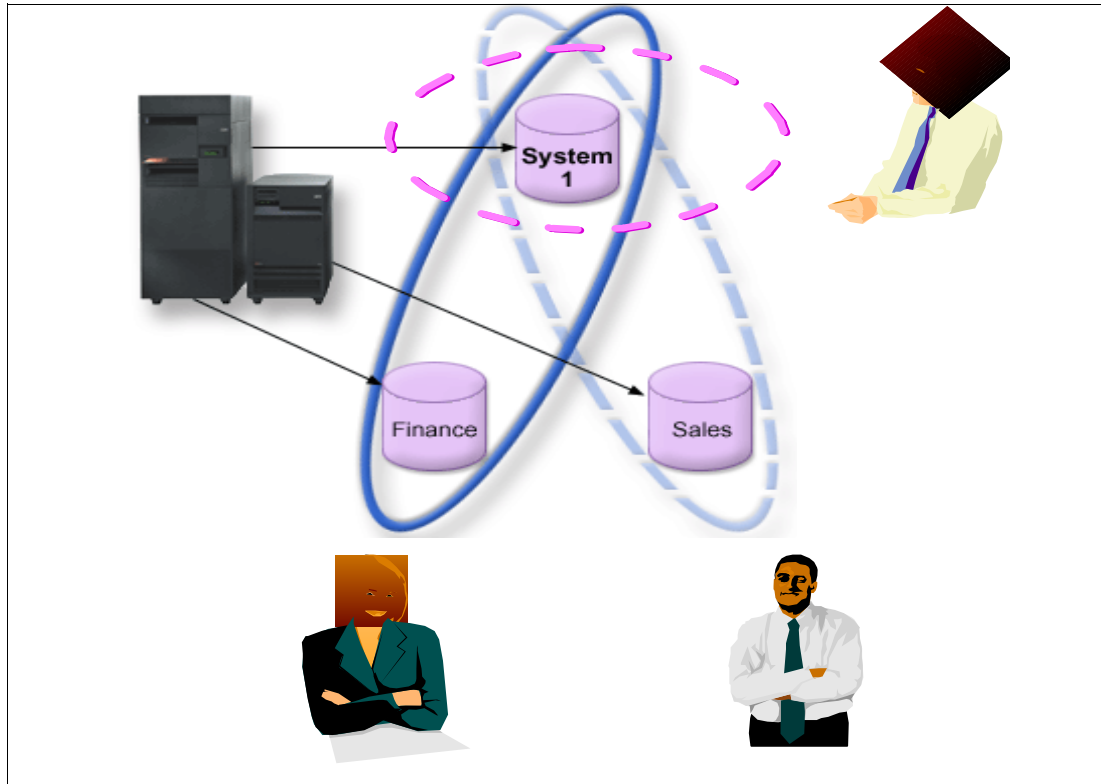


Figure 1-3 Namespace example

Restriction: Even though different ASP groups can have duplicate library names, there can be no library name in an ASP group that is also in the system ASP.

Applications and users have specific ways to change their namespace, some of these options include:

- ▶ Use the SETASPGRP command.
- ▶ In an SQL environment use the command CONNECT, to specify the correct database. To achieve the fastest performance, make sure that the database to which you perform an SQL CONNECT corresponds with Independent ASP in your current library name space. You may need to use the SETASPGRP command first to achieve this. If the SQL CONNECT function is not operating within the same library name space, the application will use Distributed Relational Database Architecture™(TM) support which can affect performance.
- ▶ Use the Change Job Description (CHGJOB) command to set the initial ASP group in the job description for a user profile.

Attention: For more information on advantages and restrictions of library namespaces in conjunction with ASP groups, see info center at:

<http://publib.boulder.ibm.com/series/v5r2/ic2924/info/rzaly/rzalyapplications.htm>

1.3.2 Multisystem clustered environment

Consider a multisystem clustered environment, where the servers are members of an iSeries cluster, and an IASP is associated with a switchable hardware group in that cluster. In this environment, IASPs can be switched between systems without having to perform an IPL. In this environment, the IASP is switchable because it resides on a switchable device. This device can be an external expansion unit or tower. Or it can be an IOP on the bus shared by two or more logical partitions.

This switch can be either automatic, in the case of an unplanned outage; or manual by administering a switchover. The independent ASP can be switched because it is self contained. This can be a significant advantage because it allows for continuous availability of data, one of the primary benefits of independent switchable ASPs.

Figure 1-4 shows an example of a switchable IASP. It has two system units and four expansion units. The lower three expansion units are owned by the system unit on the left of the diagram. The upper expansion unit is owned by the system unit on the right. The lower three expansion units are switchable IASPs and are normally used by the system on the left.

In the event of an outage, either planned or unplanned, these three expansion units can be switched to the backup system on the right. This can be done manually in the case of a planned outage (for example, in the event of a system software upgrade) or automatically by using Cluster Resource Services in the event of an unplanned outage such as a hardware failure on the system unit.

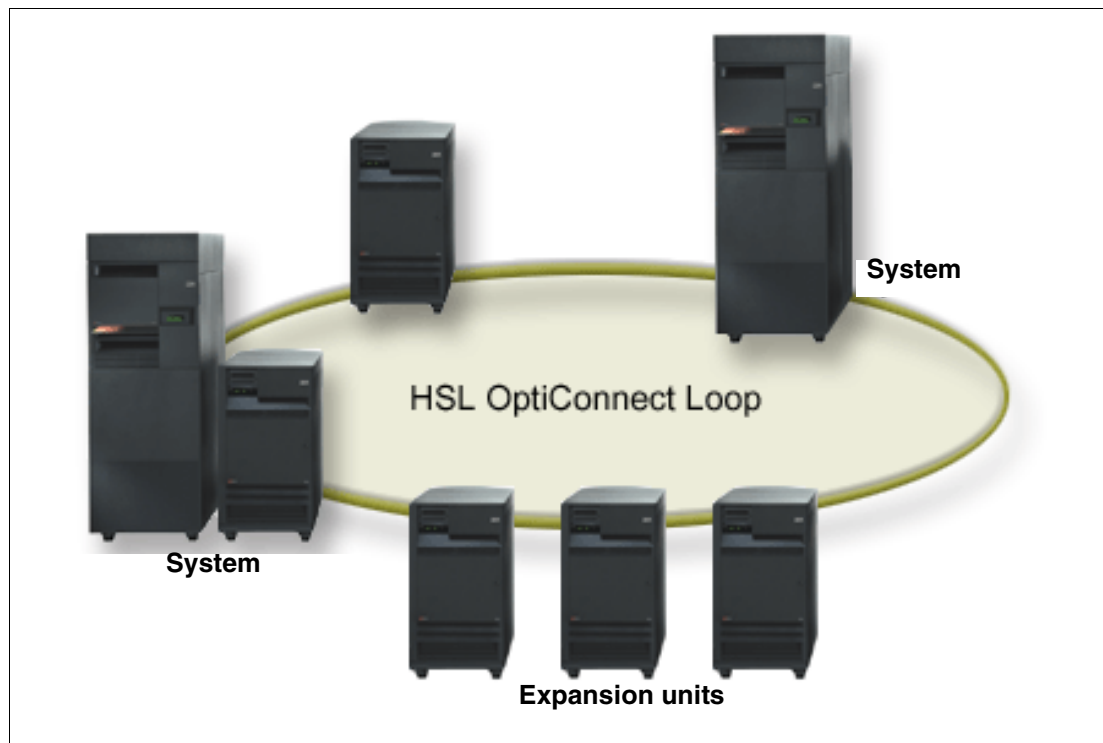


Figure 1-4 Multiple System IASP Example

Some customers are considering IASP as an alternate solution to data replication. On Figure 1-5 you can see a quick summary of how a switchable IASP solution can compare and contrast to a Data Replication solution.

	Replication	IASP
Flexibility	10's of systems	2 to 3+ systems
Single Point Of Failure	none	disk subsystem
Cost factors	Additional disk capacity. Replication software.	Switchable I/O expansion tower
Performance factors	Replication overhead	(none)
Typical failover time	5 - 30 minutes	~ 15 min
Typical switchover time	~ 5 minutes	~ 5 minutes
Real time coverage	Anything journaled	Anything on IASP
Geographic dispersion	Unlimited	Limited attach distance
Concurrent backup	Yes	No
Setup	Replication environment. What to replicate.	Device domain environment. Populate IASP.

Figure 1-5 Compare and Contrast Data Replication and IASP

Figure 1-5 IASP estimates assume near optimal system conditions for the ASP as we discuss further in this redpaper. This figure also assumes that we have few database objects in SYSBAS to track.

Important: For other information on setting up IASPs for optimal performance see the information center at:

<http://publib.boulder.ibm.com/iseres/v5r2/ic2924/info/rzaly/rzalystructure.htm>

1.4 Performance requirements

With the latest enhancements of IASP in V5R2 many customers have started moving some of their applications to IASPs, other are in the planning phase of doing it and others are in the inquiry process about the topic.

Some of the questions that customers are asking about IASP are:

- ▶ How big does SYSBAS have to be relative to the IASP(s)? What is the minimal size of SYSBAS?
- ▶ How many arms should I plan for in my IASP?
- ▶ How many arms should I plan for in SYSBAS?
- ▶ How does user profile differences affect failover/switchover time? How does the user profile synchronization affects the time?
- ▶ What kind of failover/switchover time can I expect with different configurations?

The purpose of this paper is to try to answer some of these questions.



Testing environments

In this chapter we describe the hardware and software environments that were employed to conduct the various tests described in this paper.

2.1 Cluster hardware configurations

As you can see from Figure 2-1, we had two types of environments. First we had data from an actual customer situation. Second we had two test environments which we are going to call the ITSO Lab and the Cluster Lab. The figure below shows the basics of each environment.

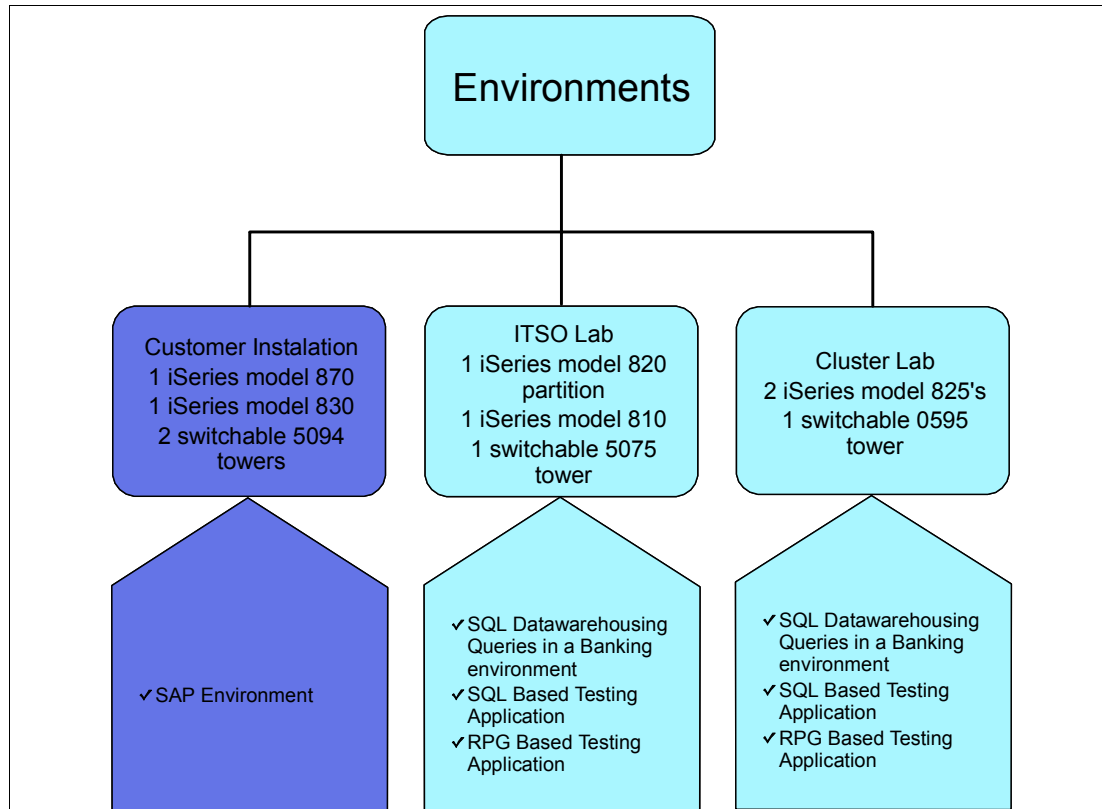


Figure 2-1 Summary of the different environments

2.1.1 Customer SAP environment

VEKA AG with headquarter in Sendenhorst (Germany) and many worldwide subsidiaries is one of the leading system developers and system providers of plastic profiles for doors, windows, shutters and plastic sheets. VEKA has been running SAP on the IBM iSeries server for many years. In Chapter 3, "SAP - Customer scenario" on page 27, we discuss an implementation of SAP on IASP.

2.1.2 Lab testing environments

In Table 2-1 we lay out the base hardware configuration of the systems that we used.

For the scenarios ran in the ITSO Lab we used two physical systems, SYSTEM01 which is an 810 and SYSTEM02 which is an 820. SYSTEM01 was not partitioned. SYSTEM02 was a secondary partition. The primary partition of SYSTEM02 is called SYSTEM0, but was not used in the scope of this red paper.

For the scenarios ran in the Cluster Lab, we used two physical systems, SYSTEM03 and SYSTEM04 which both are an 825. SYSTEM03 and SYSTEM04 were both primary partitions. The secondary partitions of SYSTEM03 and SYSTEM04 were not used within the scope of this redpaper.

Table 2-1 Lab system hardware

System name	System location	System models	DASD type	CPU type	Main storage
SYSTEM01	ITSO Lab	810	4326's 0 to 6 arms configured depending on the scenario	System processor feature : 7430 Processor feature : 2469 Interactive feature : 7430 2 processors configured	6.1 G Machine Pool 0.5GB Base 5.0GB
SYSTEM02	ITSO Lab	820	6718's 10 arms configured	System processor feature : 24BE Processor feature : 2438 Interactive feature : 1527 3 processors configured for User Profile testing except where indicated 2 processors configured for DASD testing	6.1 G Machine Pool 1.5GB Base 2.5GB
SYSTEM03	Cluster Lab	825	4326's 2 arms configured	System processor feature : 7418 Processor feature : 2473 Interactive feature : 7418 6 processors configured	23.4 G Machine Pool 1.7GB Base 20.0GB
SYSTEM04	Cluster Lab	825	4326's 2 to 14 arms configured depending on the scenario	System processor feature : 7418 Processor feature : 2473 Interactive feature : 7418 6 processors configured	23.4 G Machine Pool 1.7GB Base 20.0GB

In Table 2-2 we lay out the base hardware configuration of the IASPs that we used. SYSTEM01 and SYSTEM02 had a device CRG named BLUECRG with one IASP named BLUEIASP. SYSTEM03 and SYSTEM04 had a device CRG named CRGRED with one IASP named IA01. In both cases a physical expansion tower was used instead of a switchable IOP within LPAR.

Table 2-2 Lab IASP hardware

IASP name	Systems attached to	Hardware type	DASD type
BLUEIASP	SYSTEM01 SYSTEM02	5075 Expansion Tower	6719's Used with 0 - 6 arms configured depending on the scenario
IA01	SYSTEM03 SYSTEM04	0595 Racked Expansion Tower	4326's Used with 0 - 12 arms configured depending on the scenario

Figure 2-2 shows the hardware layout of the system SYSTEM01 (the 810) and SYSTEM02 (the 820) on the HSL loop with the 5075 expansion tower. The whole expansion tower was included in the device CRG BLUECRG with the one IASP BLUEIASP. BLUEIASP is a Primary IASP and is the only member of its ASP group, which means that there is no secondary IASP assigned. At one point in the Performance findings for “SYSBAS vs IASP configurations” on page 49, in testing an equivalent non-IASP environment, we did use the

expansion tower that was an IASP as a private tower to increase the amount of DASD in SYSBAS of the 810, named SYSTEM01.

Note: As you can note from Table 2-2, this particular study was oriented more towards the entry market of iSeries customers. This due to limitations on the different hardware configurations that we available for this study. In configurations were there is more DASD arms available the results could be very different.

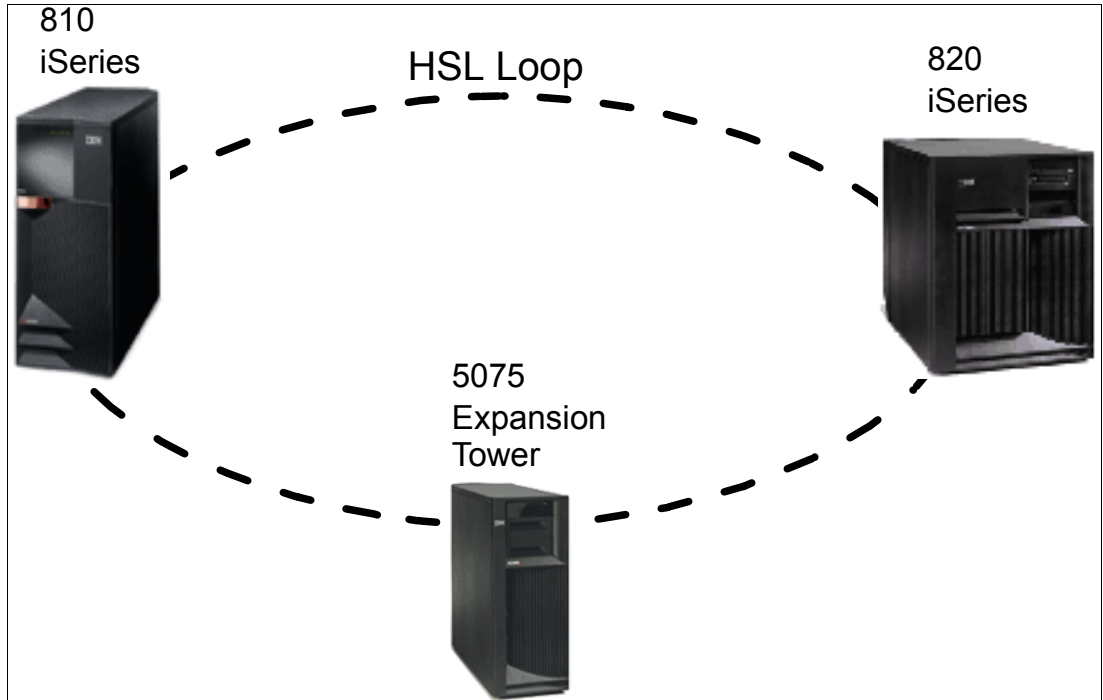


Figure 2-2 Configuration of SYSTEM01 (810) and SYSTEM02 (820) and expansion tower that served as switchable device CRG

Note: Note, this is a simplified version that should work for a basic understanding of the hardware environment that was used.

Figure 2-3 shows the output of the Display Cluster Information command (DSPCLUINF). It shows that we have the cluster Red with three nodes, SYSTEM02, SYSTEM0, and SYSTEM01. As noted before, only the nodes SYSTEM02 and SYSTEM01 were used within the scope of this red paper.

```

DSPCLUINF                               Page    1
5722SS1 V5R2M0 020719                   SYSTEM01 12/05/03 06:18:10
                                     Display Cluster Information
Cluster . . . . . : RED
Consistent information in cluster : *YES
Current cluster version . . . . . : 3
Current cluster modification level : 0
Configuration tuning level . . . . . : *NORMAL
Number of cluster nodes . . . . . : 3
Number of device domains . . . . . : 1
Detail . . . . . : *FULL
                                     Cluster Membership List
                                     Potential
                                     Node Mod Device
Node      Status      Vers Level Domain  -----Interface Addresses-----
SYSTEM02 Active        3     0 RED     1.2.34.53
SYSTEM0  Active        3     0 RED     1.2.34.24
SYSTEM01 Active        3     0 RED     1.2.34.29
                                     DSPCLUINF
                                     Page    2
5722SS1 V5R2M0 020719                   SYSTEM01 12/05/03 06:18:10
                                     Configuration and Tuning Parameters
Receive/Send heartbeat timer
ratio . . . . . : 2
Maximum retry timer ratio . . . . . : 8
Send heartbeat interval . . . . . : 3
Retry timer value . . . . . : 1
CDAT protocol timeout interval . . . . . : 2
Cluster recovery interval . . . . . : 15
Maximum retry time . . . . . : 8
Message fragment size . . . . . : 1464
Send queue overflow . . . . . : 1024
Number of bad messages threshold : 3
Number of ack messages threshold : 10
Unreachable heartbeat ack
threshold . . . . . : 1
Reachable heartbeat ack threshold: 3
Unreachable heartbeat threshold : 4
Reachable heartbeat threshold . . . . . : 4
Delayed ack timer . . . . . : 100
Message send window . . . . . : 2
Enable multicast . . . . . : 1
Performance class . . . . . : 2
Ack remote fragments . . . . . : 0
                                     DSPCLUINF
                                     Page    3
5722SS1 V5R2M0 020719                   SYSTEM01 12/05/03 06:18:10
                                     Device Domain Information
Device Domain Name  Node      Node Status
RED                 SYSTEM02  Active
RED                 SYSTEM0  Active
RED                 SYSTEM01  Active

```

Figure 2-3 Output from DSPCLUINF command giving information on the Cluster defined between SYSTEM01 and SYSTEM02

Attention: Although there are three nodes in this cluster, only two nodes were used in the test scenarios that we ran. SYSTEM0 is the primary partition of SYSTEM02 and was included in the cluster, but not used for the scope of the testing of this redpaper.

Figure 2-4 shows the output of the DSPCRGINF command. Between these cluster nodes, were defined two CRG's, REDCRG and BLUECRG. REDCRG was a switchable IOP defined between the partitions SYSTEM0 and SYSTEM02, but was not used within the scope of this redpaper. BLUECRG is a switchable 5075 expansion tower that goes between SYSTEM01 and SYSTEM02

```

DSPCRGINF                               Page    1
5722SS1 V5R2M0 020719                   SYSTEM01 12/05/03 06:18:18

                                Display CRG Information
Cluster . . . . . : RED
Cluster Resource Group . . . . . : *LIST
Consistent Information in Cluster: *YES
Number of Cluster Resource Groups: 2

                                Cluster Resource Group List
Cluster Resource Group  CRG Type  Status  Primary Node
BLUECRG                 Device  Active  SYSTEM02
REDCRG                  Device  Active  SYSTEM0
  
```

Figure 2-4 Output from DSPCRGINF command giving information on device CRG defined between SYSTEM01 and SYSTEM02

Attention: Although there are two device CRGs in this cluster, only one CRGs was used in the test scenarios that we ran. BLUECRG was the only CRG used for the scope of testing in this red paper.

Figure 2-5 shows the hardware layout of the system SYSTEM03 and SYSTEM04 on the HSL loop with the 0595 expansion tower. The whole expansion tower was included in the device CRG CRGRED with the one IASP IA01. IA01 is a Primary IASP and is the only member of its ASP group, that is there is no secondary IASP assigned.




Figure 2-5 Configuration of SYSTEM03 and SYSTEM04 and expansion tower that served as switchable device CRG

Figure 2-6 shows the output of the DSPCLUINF command. It shows that we have a cluster layout with two nodes, SYSTEM03, and SYSTEM04. The cluster RED was defined between SYSTEM03 and SYSTEM04.

```

DSPCLUINF                               Page    1
5722SS1 V5R2M0 020719                   SYSTEM04 12/05/03 06:22:14

                                Display Cluster Information
Cluster . . . . . : RED
Consistent information in cluster : *YES
Current cluster version . . . . . : 3
Current cluster modification level : 0
Configuration tuning level . . . . . : *NORMAL
Number of cluster nodes . . . . . : 2
Number of device domains . . . . . : 1
Detail . . . . . : *FULL

                                Cluster Membership List
                                Potential
                                Node Mod  Device
Node   Status  Vers Level Domain  -----Interface Addresses-----
SYSTEM03 Active   3    0 RED     1.2.56.155
SYSTEM04 Active   3    0 RED     1.2.56.160

                                DSPCLUINF                               Page    2
5722SS1 V5R2M0 020719                   SYSTEM04 12/05/03 06:22:14

                                Configuration and Tuning Parameters
Receive/Send heartbeat timer
ratio . . . . . : 2
Maximum retry timer ratio . . . . . : 8
Send heartbeat interval . . . . . : 3
Retry timer value . . . . . : 1
CDAT protocol timeout interval . . . . . : 2
Cluster recovery interval . . . . . : 15
Maximum retry time . . . . . : 8
Message fragment size . . . . . : 1464
Send queue overflow . . . . . : 1024
Number of bad messages threshold : 3
Number of ack messages threshold : 10
Unreachable heartbeat ack
threshold . . . . . : 1
Reachable heartbeat ack threshold: 3
Unreachable heartbeat threshold : 4
Reachable heartbeat threshold . . . . . : 4
Delayed ack timer . . . . . : 100
Message send window . . . . . : 2
Enable multicast . . . . . : 1
Performance class . . . . . : 2
Ack remote fragments . . . . . : 0

                                DSPCLUINF                               Page    3
5722SS1 V5R2M0 020719                   SYSTEM04 12/05/03 06:22:14

                                Device Domain Information
Device Domain Name  Node   Node Status
RED                 SYSTEM03 Active
RED                 SYSTEM04 Active

```

Figure 2-6 Output from DSPCLUINF command giving information on Cluster defined between SYSTEM03 and SYSTEM04

Figure 2-7 shows the output of the DSPCRGINF command. Both SYSTEM03 and SYSTEM04 were in the device domain RED. Between SYSTEM03 and SYSTEM04 was defined a device CRG called CRGRED. This CRG had one IASP called IA01.

```

5722SS1 V5R2M0 020719                                SYSTEM04 12/05/03 06:22:42
                Display CRG Information
Cluster . . . . . : RED
Cluster Resource Group . . . . . : *LIST
Consistent Information in Cluster: *YES
Number of Cluster Resource Groups: 1
                Cluster Resource Group List
Cluster Resource Group  CRG Type      Status      Primary Node
CRGRED                  Device      Active      SYSTEM04

```

Figure 2-7 Output from DSPCRGINF command giving information on device CRG defined between SYSTEM03 and SYSTEM04

All systems used within the redpaper were installed with CUME C3252520 and with Database group ptf #10 before testing began. At the time of the testing, these were the latest CUME and DB Fixpack available.

2.2 Application software

In this section we will describe the different applications that we used to run the different scenario testings.

2.2.1 SAP - Customer scenario

One of the application software packages that was used for the IASP performance tests was SAP R/3 on iSeries. VEKA with headquarters in Sendenhorst(Germany) and many worldwide dependencies is the leading system developers and system providers of plastic profiles for doors, windows, shutters and plastic sheets. VEKA was originally running SAP on iSeries and wanted to move to a different cluster-based solution based on Independent ASP. On Chapter 3, “SAP - Customer scenario” on page 27, we describe the results of this performance testings.

2.2.2 SQL datawarehousing queries in a banking environment

Corpbanca is a bank in Chile. In the year 2001 Corpbanca began a big project in order to change its old information systems, which were running in Open VMS, to the iSeries platform. After a large effort all the pieces of the mission critical software were finished and ready to do a “Big Bang” implementation by using an iSeries 840 server which has 8 active processors and 4 processors not used, but available via capacity on demand. The big bang took place in March of 2003 and it was a very successful start with the IBM iSeries server.

The business application software is composed of a core banking solution and an in-house developed software. The core banking solution is from IBS by Datapro in Miami, which provides a java based browser interface and a nightly batch environment written in RPG. The IBM software products used are IBM Websphere, the HTTP server powered by Apache, IBM MQ Series, and IBM CICS/400®. The database resides in DB2 libraries.

At the present time, Corpbanca has a consolidated production environment and the project is seen as a very interesting experience in the Chilean banking industry.

As a bank Corpbanca needs to generate a massive amount of daily and monthly reports for management as well for the banking bureau. Corpbanca uses SQL and stored procedures to have more flexibility and productivity. This processes are very critical and Corpbanca is always looking for a way to reduce the duration of these processes and to improve the

availability of the solution. This is one of the reasons why Corpbanca is interested in evaluating Independent ASPs as a possible solution. Some of the questions that Corpbanca wants to answer regarding IASP are:

- ▶ What is the performance of an SQL application in an IASP?
- ▶ How complicated are the application changes required when moving to an IASP?
- ▶ Can the application run when a switchover occurs?

To do this tests we selected two queries (QRY1 and QRY2) that ran in different IASP configurations. In Chapter 4, “SYSBAS vs IASP configurations” on page 49, we will the results and conclusions.

2.2.3 SQL-based benchmark testing application

Two different SQL testing script scenarios were used for the IASP performance testing. These two scenarios were:

- ▶ A single batch job that contained four different simple SQL SELECT statements and each SQL statement was executed sequentially(single threaded)
- ▶ Multiple simultaneous executing batch jobs that each contained a different simple SQL SELECT statement (multi-threaded)

All of the SQL testing script SELECT statements used would force the system to build temporary indexes prior to producing the results. The same database file was used and there were no previously established indexes in existence. To be consistent, each test system had an IPL performed and the database table restored prior to the execution of each SQL testing script scenario.

The data gathered for comparison was:

- ▶ The amount of time in seconds the SQL script took to execute as a single batch job as well as the number of seconds for all of the multiple executing SQL batch jobs took to complete
- ▶ The average percentage busy of the SYSBAS storage pool disk arm utilization
- ▶ The average percentage busy of the IASP disk arm utilization

The results are described in Chapter 4, “SYSBAS vs IASP configurations” on page 49.

2.2.4 RPG-based benchmark testing application

The same RPG testing script was used for the IASP performance tests in three different ways for the benchmark results as follows:

- ▶ RPG program
- ▶ RPG program with journaling
- ▶ RPG program with journaling and commitment control

The process for the RPG test script program would read an RPG full function database table containing the account number for a customer record. The account number column was used to chain to five subordinate/support data files sequentially and update each record that was retrieved. This process was performed again with the five subordinate/support data files with journaling enabled. Lastly, the process was performed again with not only the five subordinate/support data files having journaling enabled but also with the RPG COMMIT operation code being executed across all five subordinate/support data files at once.

The data gathered for comparison was:

- ▶ The amount of time in seconds that the test RPG script took to execute

- ▶ The average percentage busy of the SYSBAS storage pool disk arm utilization
- ▶ The average percentage busy of the IASP disk arm utilization

The results are described in Chapter 4, “SYSBAS vs IASP configurations” on page 49.

2.3 OS/400

In the following sections we will discuss how we set up the user profiles for testing on the system. We will also go through how we set up our failover and switchover scenarios. Failover and switchover testing was done with the user profiles as well as with different DASD configurations.

2.3.1 User profiles

In a clustered environment, a user profile is considered to be the same across servers if the profile names are the same. The name is the unique identifier in the cluster. However, a user profile also contains a user identification number (UID) and a group identification number (GID). To reduce the amount of internal processing that occurs during a switchover, where the independent disk pool is made unavailable (varied off) on one server and then made available (varied on) on a different server, the UID and GID values should be synchronized across the cluster domain.

When switching an IASP between systems as part of a CRG, if the UID and GID are not synchronized for each profile, the UID and GID of each object must be updated on each switch, for each object owned by the mismatched profile. Since this can take a lot of time, it is good to have the UID and GID synchronized for each profile on the systems that uses objects in the IASP.

Attention: For more information on User Profiles in a clustered environment, see info center:

<http://publib.boulder.ibm.com/iseries/v5r2/ic2924/index.htm?info/rzaly/rzalysynchronizoid.htm>

We discussed the UID/GID User Profile issue with IASP developers in Rochester and found out that the timing issue comes mostly from the number of objects that have a mismatch. The amount of user profiles that own mismatched objects, seems to be irrelevant. Also the type of object doesn't appear to matter within its file system. That is, any library object takes about the same amount of time to update as any other library object. Any IFS object takes about the same amount of time to update as any other IFS object. The size of the object also does not seem to come into play.

Knowing the above, we decided to test failover and switchover scenarios where we had UID or UID and GID mismatches for different amounts of objects. We tested the following scenarios:

- ▶ One profile with a mismatch owning 1,000 objects
- ▶ Two profiles with a mismatch owning 1,000 objects each, for a total of 2,000 objects
- ▶ Three profiles with a mismatch owning 1,000 objects each, for a total of 3,000 objects
- ▶ Four profiles with a mismatch owning 1,000 objects each, for a total of 4,000 objects
- ▶ One profiles with a mismatch owning 100,000 objects

We ran the above scenario four different times, as:

- ▶ UID mismatch for library objects

- ▶ UID and GID mismatch for library objects
- ▶ UID mismatch for IFS objects
- ▶ UID and GID mismatch for IFS objects

For library objects we used physical files. For IFS objects, we used STMFs.

Restriction: Due to time restrictions on our test systems we were not able to do UID/GID Failover or Switchover testing for library objects.

We tested on the 810 and 820 partition. See “Lab testing environments” on page 10 for more information.

Checking for profile UID/GID mismatch

Next we had to make sure that the profiles on the two systems we were testing with, were synched up for the UID and GID of the user profiles. To do this we ran the command DSPUSRPRF as shown in Figure 2-8 on both systems. Then we brought the output files to the same system. At this point we ran the SQL statement in Figure 2-9 which gave the output in Figure 2-10.

```

Display User Profile (DSPUSRPRF)

Type choices, press Enter.

User profile . . . . . > *ALL           Name, generic*, *ALL
Type of information . . . . . *BASIC     *BASIC, *ALL, *CMDAUT...
Output . . . . . > *OUTFILE           *, *PRINT, *OUTFILE
File to receive output . . . . . > FILE1     Name
Library . . . . . > LIB             Name, *LIBL, *CURLIB

Output member options:
Member to receive output . . . . . *FIRST     Name, *FIRST
Replace or add records . . . . . *REPLACE   *REPLACE, *ADD

Bottom

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 2-8 DSPUSRPRF command example, one was ran for both systems

Then we brought the output files to the same system. At this point we ran the SQL statement in Figure 2-9 which gave the output shown in Figure 2-10.

```
Enter SQL Statements

Type SQL statement, press Enter.
===> select
      a.upuprf,
      a.upsyst,
      a.upuid,
      a.upgid,
      b.upuprf,
      b.upsyst,
      b.upuid,
      b.upgid
from LIB/FILE1 a inner join LIB/FILE2 b on          s
      a.upuprf = b.upuprf
where a.upuid <> b.upuid
      or a.upgid <> b.upgid

Bottom

F3=Exit   F4=Prompt   F6=Insert line   F9=Retrieve   F10=Copy line
F12=Cancel   F13=Services   F24=More keys
```

Figure 2-9 SQL statement to check if User Profiles are Synced on UID and GID parameter for the files produced from Figure 2-8

```

Display Data
width . . . . . : 106
Position to line . . . . . _____
column . . . . .
.....1.....2.....3.....4.....5.....6.....7.....8.....9.....10.
....+
User      System      USER ID      GROUP ID      User      System      USER ID      GROUP ID
                NUMBER          NUMBER
ZZOBJ1    SYSTEM02    6,000,001    8,000,001    ZZOBJ1    SYSTEM01    5,000,001    8,000,001
ZZOBJ2    SYSTEM02    5,000,002    7,000,002    ZZOBJ2    SYSTEM01    5,000,002    8,000,002
***** End of data *****

Bottom
F3=Exit      F12=Cancel      F19=Left      F20=Right      F21=Split      F22=Width 80

```

Figure 2-10 Sample output from the SQL statement in Figure 2-9

This process worked very well to check that in each step of the process, only the profile(s) that we wanted to mismatch on the UID and/or GID were mismatched.

Tip: This process is explained in more detail in the registered Knowledge Base Document 29675522. The Registered Knowledge Base documents can be reached at:
<http://www-912.ibm.com/supporthome.nsf/document/1000057>

Tip: Management Central can also assist with synching UID's and GID's for user profiles. See the Info Center article at:
<http://publib.boulder.ibm.com/iserics/v5r1/ic2924/info/rzaih/rzaihsynchuidgid.htm>

2.3.2 Failover times

A failover occurs when a server in a cluster automatically changes over to one or more backup servers in the event of a system failure. A switchover and a failover function identically once they have been triggered. The only difference is how the event is triggered.

When a failover occurs, access is switched from the cluster node currently acting as the primary node in the recovery domain of the cluster resource group to the cluster node designated as the first backup.

Attention: For more information on Cluster Failovers, see info center:

<http://publib.boulder.ibm.com/series/v5r2/ic2924/info/rzaig/rzaigconceptsfailover.htm>

We considered several possibilities to trigger the failover scenario:

- ▶ Control panel option 8
- ▶ Control panel option 3
- ▶ PWRDWNSYS OPTION(*IMMED) command

We decided that the PWRDWNSYS (option 3) would be the best way to go to achieve consistent measurements. That is because with option 1 and 2, we could very likely get a cluster partition instead of a failure, and then it would require manual intervention to tell the cluster to fail over. With this manual intervention, we could not keep the time the failover time we desired a constant. With option 3, cluster always noted the failure instead of a partition.

We also synchronized the times between the two systems in order to obtain consistent timings.

Attention: For more information on cluster failure vs. cluster partition, see info center:

<http://publib.boulder.ibm.com/series/v5r2/ic2924/index.htm?info/rzaig/rzaigtroubleshootexamplefailover.htm>

To get the times involved in the switch we checked the history log for message CPF0930 (see Example 2-1) for the IASP on the system the IASP was moving from. On the system the IASP was moving to, we checked the history log for message CPC2605 (see Example 2-2) for that IASP.

*Example 2-1 Message CPF0930 displayed via DSPLOG OUTPUT(*PRTWRAP)*

```
CPF0930 60 INFO      PWRDWNSYS command in progress.
```

*Example 2-2 Message CPC2605 displayed via DSPLOG OUTPUT(*PRTWRAP)*

```
CPC2605 00 COMPLETION Message . . . . : Vary on completed for device BLUEIASP.  
Cause . . . . . : The vary configuration (VRYCFG) command was issued requesting that device BLUEIASP  
be varied on. Technical description . . . . . : The status of device BLUEIASP will be VARY  
ON PENDING until it is contacted or a device failure occurs.
```

2.3.3 Switchover times

A switchover happens when you manually switch access to a resource from one server to another. You would usually start a manual switchover if you wanted to perform system maintenance, such as applying PTFs, installing a new release, or upgrading your system.

When a switchover occurs, access is switched from the cluster node currently acting as the primary node in the recovery domain of the cluster resource group to the cluster node designated as the first backup.

Attention: For more information on Cluster Switchover, see info center:

<http://publib.boulder.ibm.com/series/v5r2/ic2924/index.htm?info/rzaig/rzaigconceptsswitchover.htm>

To do the switchover testing we used the command CHGCRGPRI (see Figure 2-11) and supplied the Cluster and CRG involved in the switch.

```

Change CRG Primary (CHGCRGPRI)

Type choices, press Enter.

Cluster . . . . . RED      Name
Cluster resource group . . . . BLUECRG  Name
Exit program data . . . . . *SAME

_____
_____
_____

Bottom
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

```

Figure 2-11 Sample CRGCRGPRI command used in our testing to initiate a switchover.

To get the times involved in the switch we checked the history log for message CPC2606 (see Example 2-3) for the IASP on the system the IASP was moving from. On the system the IASP was moving to, we checked the history log for message CPC2605 (see Example 2-2 on page 22) for that IASP.

*Example 2-3 Message CPC2606 displayed via DSPLOG OUTPUT(*PRTWRAP)*

```

CPC2606 00 COMPLETION Message . . . . : Vary off completed for device BLUEIASP.
Cause . . . . : The vary configuration (VRYCFG) command was issued requesting that device NON_SWITCH
be varied off. Technical description . . . . . : If ASCVRYOFF(*YES) was specified on the
VRYCFG command, the status of device NON_SWITCH will be VARY OFF PENDING until device take down is
completed. If ASCVRYOFF(*NO) was specified on the VRYCFG command, device take down is complete and the
status of device NON_SWITCH is VARIED OFF.

```



Part 2

Performance findings

In this part we will illustrate the different performance findings such as:

- ▶ A real customer scenario using SAP
- ▶ Different SYSBAS vs IASP DASD configurations
- ▶ How user profile synchronization affects the switchover times



SAP - Customer scenario

This section describes how SAP R/3 on iSeries performs in a switchable IASP cluster environment. All data is based on the experiences we made with the installation of customer VEKA AG in Germany. We describe the original configuration and later we discuss the changes to this configuration to use IASPs. We will also discuss our findings.

3.1 Introduction - The company profile VEKA AG

VEKA AG with headquarter in Sendenhorst (Germany) and many worldwide subsidiaries is one of the leading system developer and system provider of plastic profiles for doors, windows, shutters and plastic sheets. A wide product range, high quality standard and a solid company policy account for the success of VEKA AG not to mention an unmistakable company profile.

Since 1969 VEKA has developed into an industry-leading international group of companies. During this time, the capacities of VEKA AG have also been continuously developed. Today, VEKA has its own, high precision mixing plants in which virgin PVC and corresponding additives are mixed according to specific formulations and directly processed on the most modern extrusion machines. With the range of profile geometries offered by VEKA AG, every type of window desirable can be manufactured.

3.2 VEKA's solution overview

This section shows how VEKA originally was running SAP on iSeries and why VEKA wanted to change this landscape. On addition we will introduce the new cluster solution and the advantages the customer is expecting from this.

3.2.1 SAP architecture overview

An SAP system has one database server and one or more application servers. The application server contains the executional environment which is a collection of processes called SAP instance. An instance is connected to only one database, multiple instances can be defined for a single SAP system. An SAP database is implemented as an SQL-collection including commitment control functionality with the library name R3<SID>DATA (In our case R3VAPDATA), where <SID> is the name of the SAP-system and an SAP instance is implemented as an iSeries subsystem with SAP work processes running in it. Work processes are jobs within the instance subsystem that actually perform the work.

3.2.2 The original installation

VEKA has been running SAP on IBM iSeries for many years in a so-called SAP 2-tier environment which means that the application and database layer of the SAP solution suite is running on one central system (in this case an iSeries i830 8-way machine with a LPAR for SAP R/3 consisting of 5 processors, 12 GB Memory and 446 GB disk storage (internal 10k RPM disk technology). All data is stored in the system ASP(ASP1), apart of the journal receivers which are stored in another ASP(ASP2) as shown in Figure 3-1.

VEKA is running OS/400 V5R2 and SAP database Release 4.6C with SAP kernel release 4.6D.

Around 350 users are connected to the system using a wide mix of different SAP modules inside the SAP R/3 system

VEKA's SAP Landscape on a single system

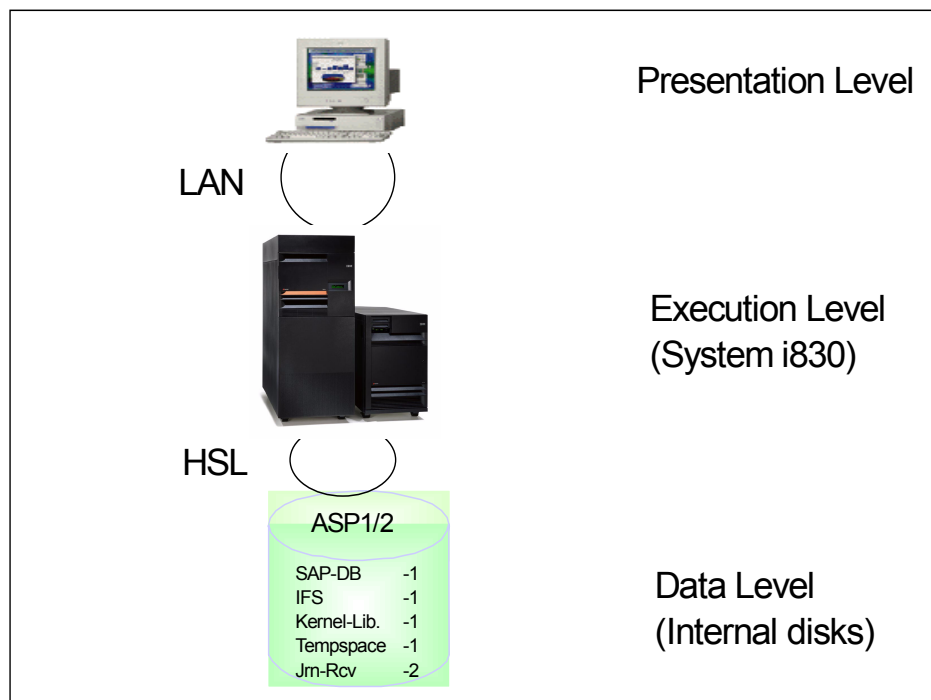


Figure 3-1 VEKA's 2-tier landscape, non-clustered environment

Figure 3-2 shows the SAP implementation in more detail. Apart of the journal receiver library (R3VAPJRN) all data of the SAP environment is stored in the ASP1 such as the SAP database, libraries containing the SAP kernel, work management objects, SQL-packages, the SAP directory structure in the IFS and the temporary storage.

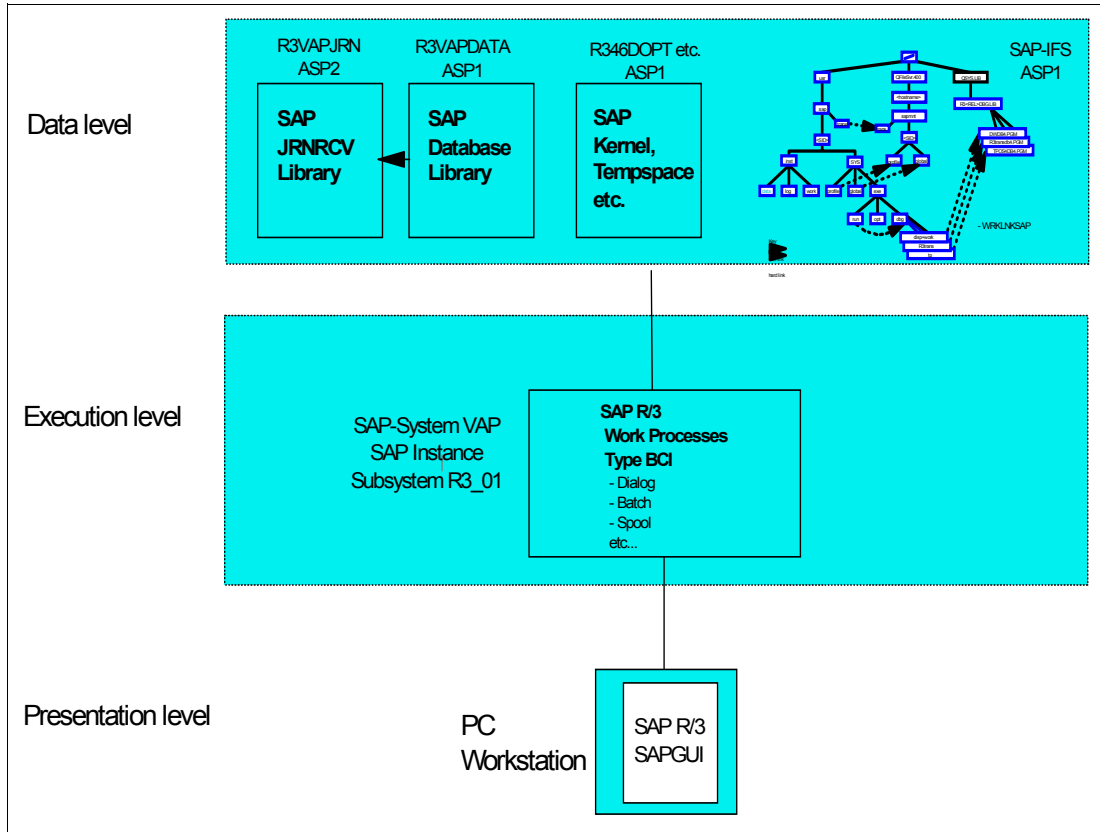


Figure 3-2 VEKA's original SAP implementation on iSeries

For Backup and High Availability reasons a solution for permanent data replication from a third party vendor was installed which permanently synchronized a duplicate of the productive database on an i830 backup server as shown in Figure 3-3.

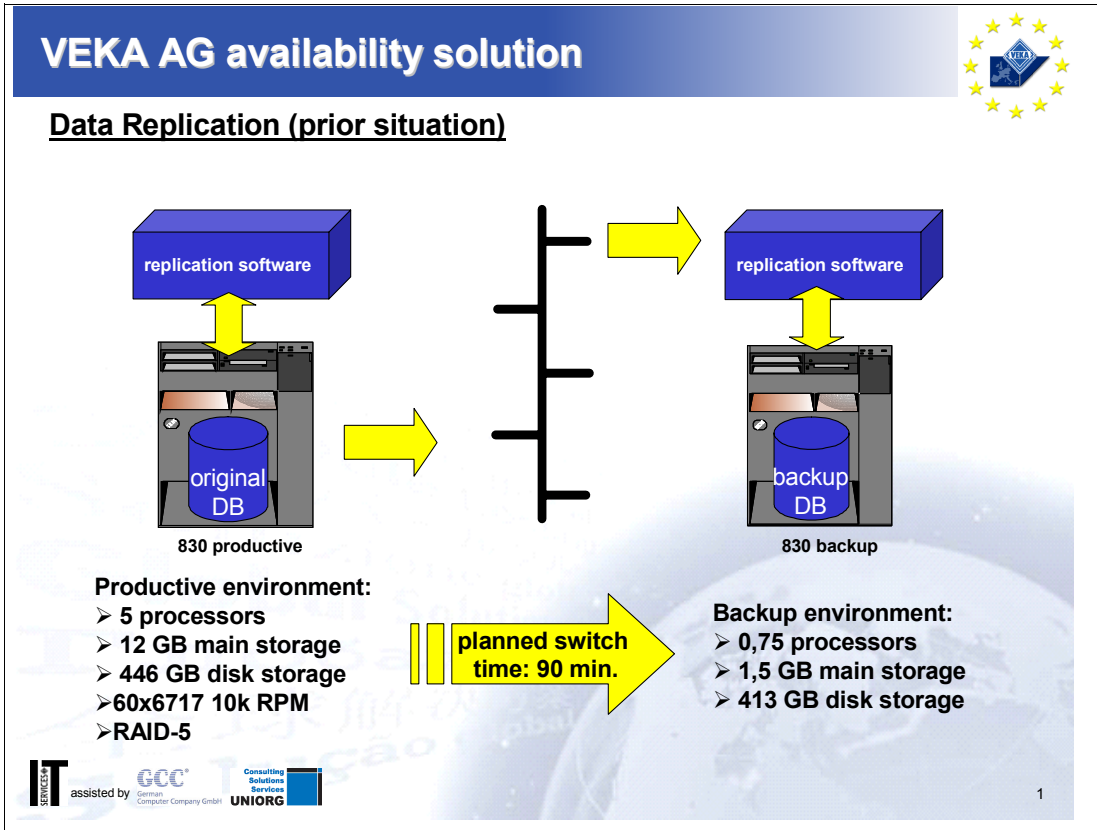


Figure 3-3 VEKA's original configuration

For many reasons VEKA was not totally satisfied with this solution:

- ▶ Switchover procedure of about 90 minutes was too long
- ▶ High complexity of switchover and apply process
- ▶ Too many partners being involved in case of support (Application provider, IBM, ISV for data replication SW)
- ▶ No disaster recovery - all servers were placed in one Operating Center

3.2.3 The new cluster-based solution design

With the enhancements of IASP in V5R2, such as the support of SAP R/3 on switchable Independent Auxiliary Storage Pools (IASP's) and VEKA's plans to build a second Operating Center for backup reasons with the Optical HSL range of 250 meters caught the eye of VEKA's technical staff as a medium cost HA-alternative that could also provide limited Disaster/Recovery(D/R) by allowing separate copies of data to reside in separate buildings up to 250 meters apart. VEKA and the IBM technical representative worked out a proposal for a new solution which changed the SAP environment into a switchable IASP landscape, as shown Figure 3-4.

VEKA AG availability solution

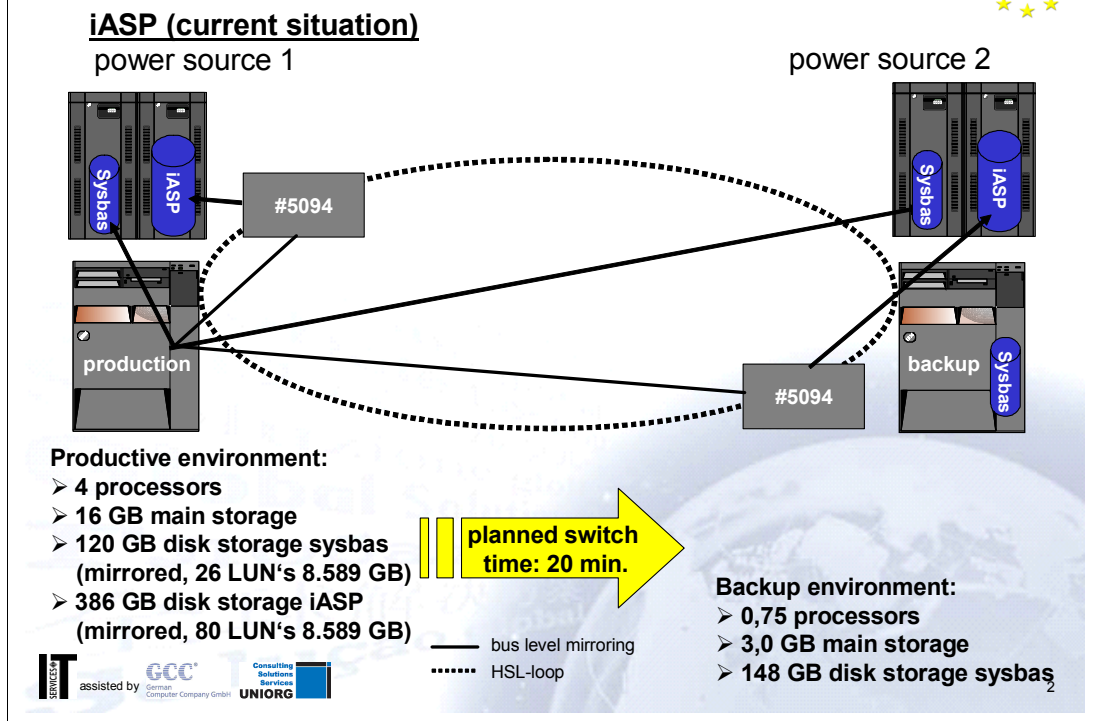


Figure 3-4 VEKA - new solution overview

In this new environment SAP R/3 will run on an i870 in a 4-way LPAR with non-switchable attached disks (SYSBAS) and a switchable IASP which contains the SAP database and all other parts of SAP which can be changed dynamically. All disks are mirrored (bus level protection) which means that one copy is placed both on the Enterprise Storage Server® (ESS) in Operating Center 1 and another one in Operating Center 2. For planned or even unplanned downtimes of the production system the IASP disk environment (two switchable towers #5094) can be switched over to the backup system being placed in the new backup Operating Center 2. This even allowed a certain level of disaster recovery. In case that one Operating Center completely falls out there is still one SAP server with a local SYSBAS and one copy of the switchable IASP available.

As long as the productive i870 is up and running the backup environment is also available, but only with reduced resources (0,75 processors, 1,5 GB main memory). In case of switchover the customer will reallocate the LPAR's and extend the hardware resources of the backup environment.

With the help of IBM's Benchmark Center in Montpellier (France) and the iSeries Technology Center in Rochester, IBM provided the proposed hardware and IASP skills in order to:

- ▶ Install the VEKA environment into a switchable IASP structure.
- ▶ Create an iSeries cluster and make the IASP switchable.
- ▶ Perform regular switchover tests.
- ▶ Perform failovers to prove the data and SAP system will move to the running system and be available to run production work.

By the end of the four days VEKA's technical personnel had successfully completed all of these tasks and were convinced that the proposed solution would fulfill their requirements:

- ▶ Switchover time of about 15 to 20 minutes
- ▶ Low complexity of switchover process
- ▶ Reduced complexity in case of support because only two partners (Application Provider and IBM) being involved
- ▶ Improved disaster recovery concept because of two operating centers.

Based on these facts VEKA finally decided for the proposed solution with a production and a backup system splitted up into different LPAR's:

- ▶ i870 into LPAR's for SAP R/3, SAP BW and Lotus® Domino
- ▶ i830 into LPAR's for SAP R/3 Backup and Test/Development environment.

We only want to focus on the SAP R/3 production and the R/3 Backup environment, which have dedicated expansion units for the SYSBAS disks and the IASP disks. All data (both SYSBAS and IASP) are placed in 2 IBM ESS storage servers Mod. 800, one with 10k RPM, the other one with 15k RPM disk technology.

3.3 SAP installation overview

This section shows how the VEKA SAP R/3 system was installed in a switchable IASP landscape. VEKA continues to run SAP R/3 in a 2-tier environment as explained in "The original installation" on page 28. The data now is split into two parts one which contains static data which normally will not be changed at execution time and which is still located in the SYSBAS and another part located in the switchable IASP which contains data being changed dynamically by the application. The SYSBAS will be non-switchable and physically available on both machines while the Switchable IASP group (ASPVAPDATA) is single and can be assigned to both machines (but only to one at a time) as shown in Figure 3-5.

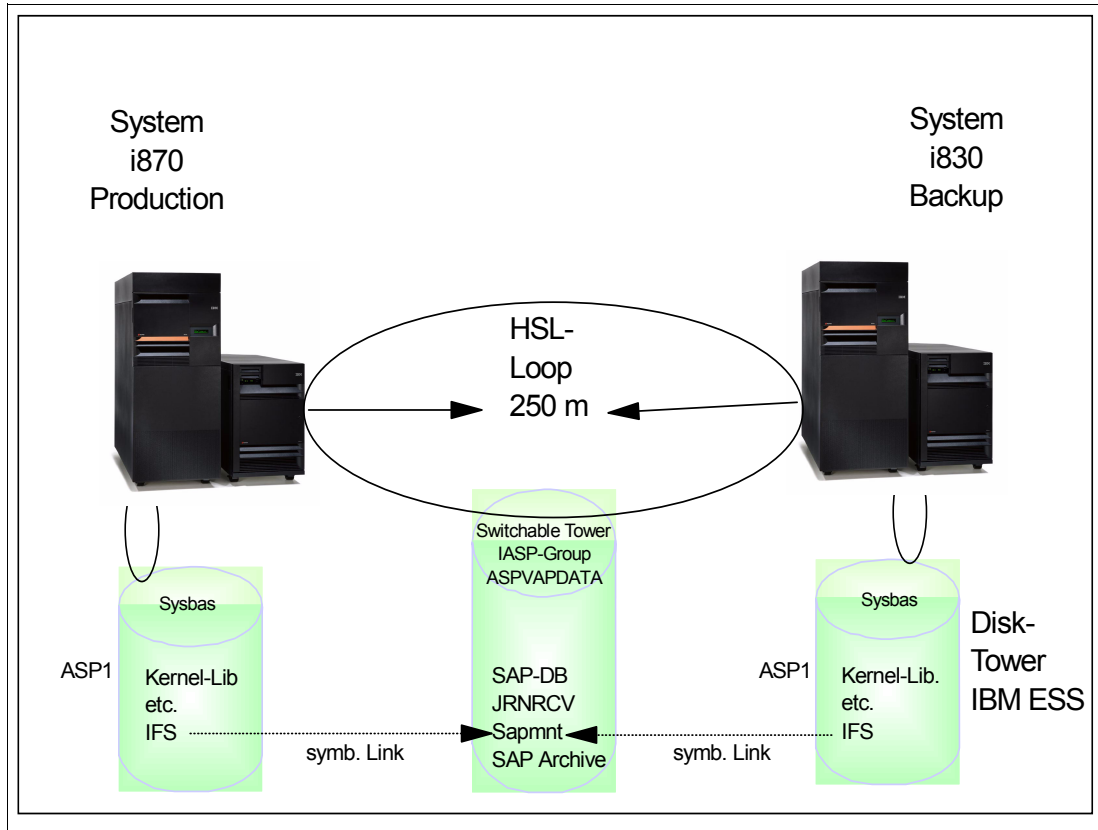


Figure 3-5 SAP landscape for Switchable IASP

Figure 3-6 shows in more details how the data was split up.

SYSBAS contains:

- ▶ The OS/400 environment
- ▶ The SAP kernel and work management libraries
- ▶ The libraries containing the SQL packages
- ▶ The temporary storage of the execution environment
- ▶ Parts of the SAP file system structure

The switchable IASP group ASPVAPDATA contains:

- ▶ The SAP database (R3VAPDATA)
- ▶ Parts of the SAP file system structure (sapmnt-directory, SAPArchive-directories)
- ▶ The journal receiver library (secondary IASP)

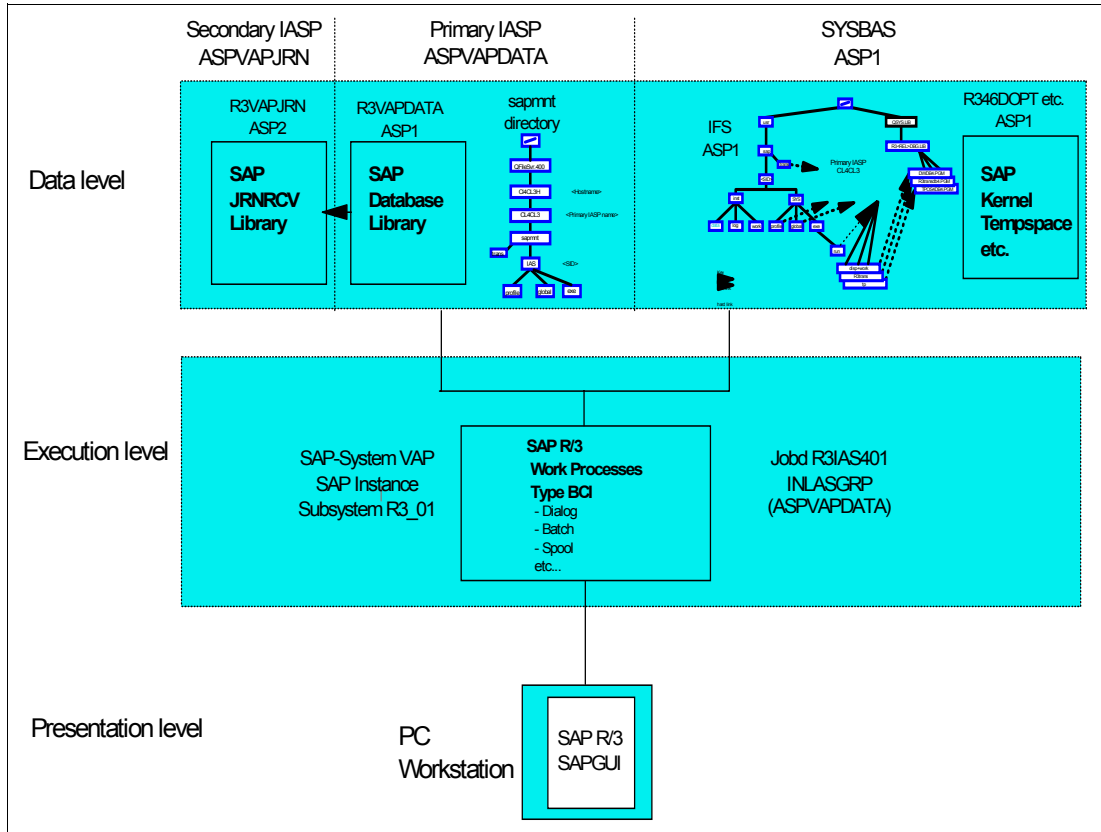


Figure 3-6 SAP implementation in IASP

For more information about how to implement an SAP R/3 system into an Independent ASP refer to Chapter 10 of *IBM eServer iSeries Independent ASPs: A Guide to Moving Applications to IASPs*, SG24-6802.

3.4 The cluster environment

In this section we illustrate how the cluster environment for supporting the SAP solution in an IASP environment was defined. In Figure 3-7 we illustrate the SAP running in a cluster environment.

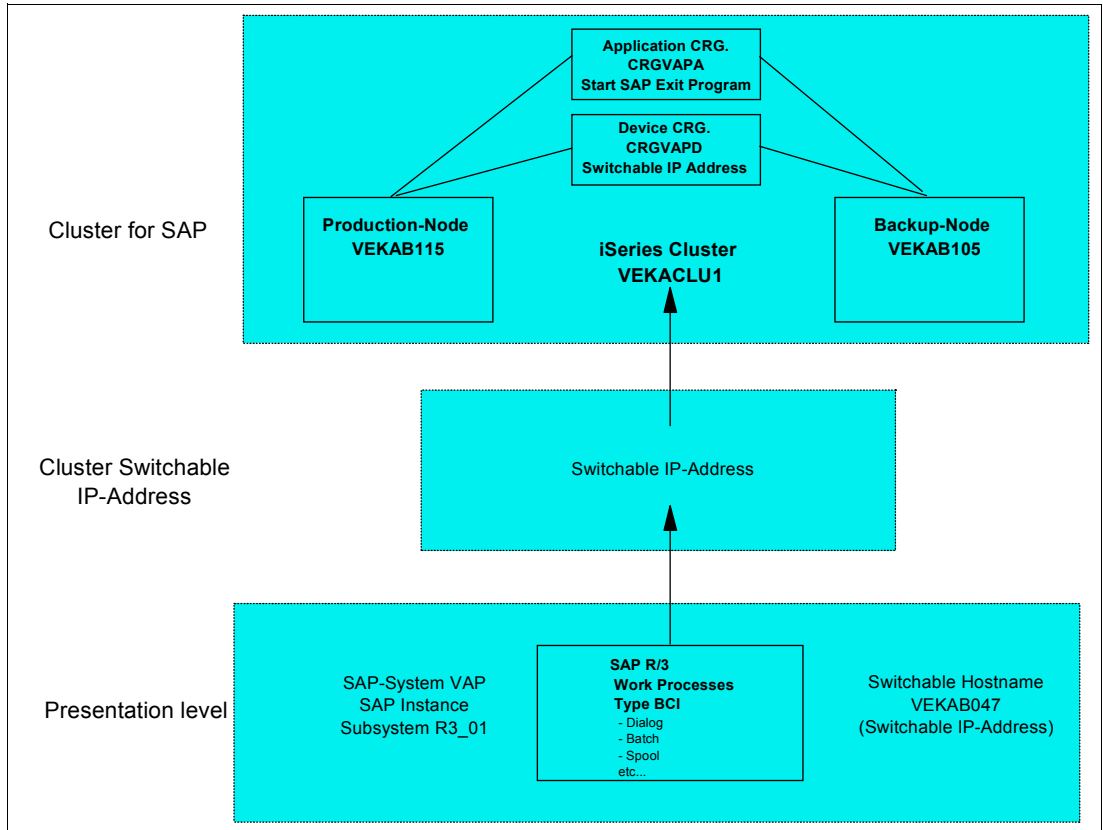


Figure 3-7 SAP running in VEKA-cluster VEKAVCLU1

In Figure 3-8 illustrates the Display Cluster Information for the cluster defined at VEKA. It has defined a Production node (VEKAB115) and a Back-up node (VEKAB105).

```

Display Cluster Information

Cluster . . . . . : VEKACLU1
Consistent information in cluster : *YES
Current cluster version . . . . . : 3
Current cluster modification level : 0
Configuration tuning level . . . . : *NORMAL
Number of cluster nodes . . . . . : 2
Number of device domains . . . . . : 1
Detail . . . . . : *FULL

                Cluster Membership List

                Potential
                Node Mod  Device
Node   Status  Vers Level Domain  -----Interface Addresses-----
VEKAB115 Active    3    0 VEKACLU1  99.999.9.115  99.999.9.116
VEKAB105 Active    3    0 VEKACLU1  99.999.9.105  99.999.9.106

                                                                Bottom

Enter=Continue  F12=Cancel  F3=Exit  F1=Help

```

Figure 3-8 Display Cluster Information

Figure 3-9 illustrates the Display CRG Information. It shows the cluster resource groups such as the Application CRG (CRGVAPA) and also the Device CRG (CRGVAPD).

```

Display CRG Information

Cluster . . . . . : VEKACLU1
Cluster Resource Group . . . . . : *LIST
Consistent Information in Cluster: *YES
Number of Cluster Resource Groups: 2

                Cluster Resource Group List

Cluster Resource Group  CRG Type  Status  Primary Node
CRGVAPA                 Application Active   VEKAB115
CRGVAPD                 Device    Active   VEKAB115

                                                                Bottom

Enter=Continue  F12=Cancel  F3=Exit  F1=Help

```

Figure 3-9 VEKA Cluster Resource groups

3.4.1 The device CRG

The device CRG ASPVAPD enables the hardware resources (in our case the primary IASP ASPVAPDATA and the adjacent secondary one ASPVAPJRN) to be switched between our systems i870 and i830. In addition, the IP takeover function allows to take over the IP address of the primary system in case of a switchover as you can see in Figure 3-10.

```
Display CRG Information

Cluster . . . . . : VEKACLU1
Cluster Resource Group . . . . . : CRGVAPD
Reporting Node Identifier . . . : VEKAB115
Consistent Information in Cluster: *NO

                                Configuration Object Information

Configuration Object Device Device Vary Server
Object Name Type Type Subtype Online Ip Address
ASPVAPDATA *DEVD *ASP Primary *YES 99.999.9.1
ASPVAPJRN *DEVD *ASP Secondary *PRIMARY *NONE

Number of Device List Entries . . : 2

Enter=Continue F12=Cancel F3=Exit F1=Help

Bottom
```

Figure 3-10 The device CRG

In Figure 3-11 you can see the detailed information regarding the device CRG.

```

Display CRG Information

Cluster . . . . . : VEKACLU1
Cluster Resource Group . . . . . : CRGVAPD
Reporting Node Identifier . . . . . : VEKAB115
Consistent Information in Cluster: *YES

Cluster Resource Group Type . . . . . : Device
Cluster Resource Group Status . . . . . : Active
Previous CRG Status . . . . . : Switchover Pending
Exit Program . . . . . : *NONE
  Library . . . . . : *NONE
Exit Program Format . . . . . : *NONE
Exit Program Data . . . . . : *NONE

User Profile . . . . . : *NONE
Text . . . . . :

Enter=Continue  F12=Cancel  F3=Exit  F1=Help

More...

```

Figure 3-11 Device CRG - Detailed description

3.4.2 The application CRG

The application CRG CRGVAPA enables the application program R3IASPEXIT to be restarted. The application program R3IASPEXIT (an example is delivered by SAP) starts the SAP system whenever the application CRG will be activated. This allows to start the SAP R/3 system automatically whenever a switchover has been done. Figure 3-12 and Figure 3-13 illustrates shows the definition of the application CRG.

```

Display CRG Information

Cluster . . . . . : VEKACLU1
Cluster Resource Group . . . . . : CRGVAPA
Reporting Node Identifier . . . . . : VEKAB115
Consistent Information in Cluster: *YES

Cluster Resource Group Type . . . : Application
Cluster Resource Group Status . . : Active
Previous CRG Status . . . . . : Switchover Pending
Exit Program . . . . . : R3IASPEXIT
  Library . . . . . : SAPIASP
Exit Program Format . . . . . : EXTP0100
Exit Program Data . . . . . : VAP01ASPVAPDATA
  SLIC_PRI

User Profile . . . . . : CRGVAPA
Text . . . . . :

Job Name . . . . . : R3VAP01

Enter=Continue  F12=Cancel  F3=Exit  F1=Help
More...

```

Figure 3-12 Application CRG - Detailed description

```

Display CRG Information

Cluster . . . . . : VEKACLU1
Cluster Resource Group . . . . . : CRGVAPA
Reporting Node Identifier . . . . . : VEKAB115
Consistent Information in Cluster: *YES

Distribute Information Queue . . . : *NONE
  Library . . . . . : *NONE
Failover Message Queue . . . . . : *NONE
  Library . . . . . : *NONE
Takeover IP Address . . . . . : 99.999.9.100
Allow Application Restart . . . . . : No
Number of Restarts . . . . . : 0
Configure Takeover IP Address . . : User
CRG Extended Attribute . . . . . : *NONE

Enter=Continue  F12=Cancel  F3=Exit  F1=Help
Bottom

```

Figure 3-13 VEKA - Application CRG, detailed description

3.5 Conclusions and findings

In this section we will describe the conclusions from the tests done with VEKA.

3.5.1 Failover/switchover times

We tested different scenarios in order to figure out how long a switchover of the switchable IASP would take. The scenarios tested were:

- ▶ Planned and unplanned switchovers
 - Planned switchover (Ex. for maintenance of the production system) and the way back. In this case we ordinary stop SAP, use the CHGCRGPRI command for switchover and started the device and application CRG on the target system which brought up SAP automatically.
 - Unplanned switchover (caused by a 'crash' of the production environment). In our case we did this by powering down the machine while SAP was up and running).
- ▶ Both directions, which means both a switchover from i870 production system to the i830 backup machine and the way back
- ▶ We also tested different variations from the proposed configuration in order to check if this would influence the switchover time:
 - Using only one ESS storage server with RAID-1 protection being switched off
 - Using two ESS storage server with RAID-1 (bus level protection) being enabled.
 - trying three different configurations of the 830 backup LPAR:
 - "Small" with 0,75 processors and 1,5 GB memory
 - "Large" with 1,75 processors and 5,5 GB memory
 - "Very large" with 2.75 processors and 8 GB memory

Note that we did not vary the amount of switchable towers. Even in the scenario with only one ESS we switched the corresponding expansion tower as a part of the device CRG without any disks being assigned.

The results are shown in Table 3-1.

Table 3-1 Switch test RAID-1 (2 ESS) and RAID-5 (1 ESS only)

	PWRDWN SYS ('Unplanned')		CHGCRGPRI ('Planned') (1. Device, 2. Application CRG)	
	Disk ratio: 12/36	Disk ratio: 13/42	Disk ratio: 12/36	Disk ratio: 13/42
"Small" env. 830, 2 ESS 0,75 proc., 1,5 GB main stg.				
870 >> 830	21 min.	22 min.	Not tested	31 min. (29+2)
830 >> 870	11 min.	21 min.	Not tested	18 min. (15+3)
"Small" env. 830, 1 ESS 0,75 proc., 1,5 GB main stg.				
870 >> 830	Not tested	Not tested	20 min. (17+3)	Not tested
830 >> 870	Not tested	Not tested	15 min. (12+3)	Not tested
"Large" env. 830, 2 ESS 1,75 proc., 5,5 GB main stg.				
870 >> 830	19 min.	20 min.	Not tested	23 min. (20+3)

830 >> 870	15 min.	16 min.	not tested	18 min. (15+3)
"Large" env. 830, 1 ESS 1,75 proc., 5,5 GB main stg.				
870 >> 830	16 min.	Not tested	16 min. (15+1)	Not tested
830 >> 870	11 min.	Not tested	11 min. (10+1)	Not testet
"Very large" env. 830, 2 ESS 2,75 proc., 8,0 GB main stg.				
870 >> 830	Not tested	22 min.	Not tested	Not tested
830 >> 870	Not tested	16 min.	Not tested	Not tested

Conclusions

Apart of one of the measurements the switchover time was between 11 and 23 minutes. From our testing results as documented in Table 3-1 you easily can derive two conclusions for our scenario. The switchover time:

- ▶ Takes longer if you are switching from i870 to i830 than doing the other way from i830 to i870
- ▶ Does not strongly depend on other hardware variations like giving more processors or more memory to the i830 LPAR

Summarizing you can say that in our scenario the switchover time primarily is influenced by the processor power of the target system where you want to switch your towers.

Note that in our tests we didn't see a lot of database recovery. This depends on the size of open commit cycles which need to be rolled back and may be different in other SAP-environments.

3.5.2 I/O performance considerations

We monitored the I/O performance of the new disk layout for the i870 production environment with switchable IASP and SYSBAS and compared these values with the results of the former i830 production system (non-switchable system ASP only). Three measurements being representative for the VEKA production environment were used for this comparison:

- ▶ OLD: One measurement for the original environment on i830
- ▶ Two measurements for the i870 with switchable IASP:
 - NEW1: Workload comparable to OLD
 - NEW2: High workload (100% workload increase comparing with OLD)

The results are shown in Table 3-2. All data being shown there is derived either from OS/400 Performance Reports (System Report or Resource Report). For configuration details of the old and new landscape refer to Figure 3-3 on page 31 and Figure 3-4 on page 32.

Table 3-2 I/O performance comparison.

	OLD i830 5 Way. 1756 SAPS capacity	NEW1 i870 4-Way 2750 SAPS capacity	NEW2 i870 4-Way 2750 SAPS capacity
CPU %	28.2	22	36.3
Actual Workload (SAPS)	495	605	998

	OLD i830 5 Way. 1756 SAPS capacity	NEW1 i870 4-Way 2750 SAPS capacity			NEW2 i870 4-Way 2750 SAPS capacity		
Faults per sec Non-DB/DB Pages/sec	152 322/899	95 224/702			114 192/960		
		Total	SYSBAS	IASP	Total	SYSBAS	IASP
IO/sec. (read/write ratio)	733 (582/151)	631 (465/166)	184	447	932 (748/183)	182	750
K per IO	14.8	9.6	12.5	7	7.4	13	6
K per sec.	10,484	6,057			6,896		
Disk Arm Utilization %	5	1.9	1.9	1.9	2.3	2.3	2.3
Disk Response Time/ms Service/Wait Time	8 4/4	5.5	3 2.5/0.5	6.5 3.5/3	4.9	3 2.5/0.5	5.6 2.8/2.8
IOP Utilization %	2.3	2.9			2.5		
DSK CPU Utilization %	10	18.5/15.5			18.5/15.5		
Percent Full %	75		48.3	79.5		53.7	79.9

The SAP system was moved from an i830 5-way LPAR to an i870 4-way LPAR. This is a more powerful environment - calculated in SAPS (SAP Application Performance Specification - an SAP defined measurement for CPU throughput) we moved from a 1756 SAPS environment to a 2750 SAPS environment (both values are based on the IBM sizing tables for SAP Release 4.6C which customer VEKA is running currently). This fact is reflected by the CPU-load which only increased in 30% although the total workload increased in 100% comparing to the old environment.

We increased the memory from 12GB (i830) to 16 GB (i870) which reduced the amount of memory faults per second although the workload was increasing.

We also changed the disk interface from internal disks (60 units with 10k RPM technology based on RAID-5 protection) to external disks (for SYSBAS 26 LUNs, for switchable IASP 80 LUNs, both split up to both ESS systems because of RAID-1 bus level protection). Both ESS Mod. 800 systems are running with 35 GB disk units, one of them with 10k RPM, the other one with 15k RPM disk technology.

Considering these facts Table 3-2 shows interesting facts about the I/O behavior of a switchable IASP environment in comparison with a 'traditional' one based on non-switchable system ASP.

Usage of disk space

The data of the whole environment is split up into SYSBAS and IASP.

- ▶ In our example the size of SYSBAS seems to be appropriate. Around 50% of the disk space is used, which is a usage of about 55 GB. 40 GB of that value is consumed by temporary storage, the remaining part by OS/400, SAP kernel library, other SAP work management libraries, SQL package libraries and IFS flat files.
- ▶ the IASP contains the SAP-DB, the SAP sapsmnt-directory files and additional directories containing the SAPArchive data. Today 80% of the available space is filled up with data. The customer will extend this space very soon.

I/O per second

The amount of I/Os per second is primarily depending on the application workload profile and the hardware environment where the application is running. In our measurements the workload profile remained the same. We only varied the hardware configuration between scenario 'OLD' and the NEW1 and NEW2 one. We also extended the throughput between NEW1 and NEW2 while the hardware environment was not changed.

Effect of hardware change

We changed to more powerful processor and increased the memory from 12 to 16 GB. For this reason the amount of faults per seconds were reduced for a similar workload. You can build a ratio between I/O per second and SAPS (measurement of throughput) and will find the following values:

- ▶ OLD: 1.48 IO/SAPS
- ▶ NEW1: 1.04 IO/SAPS
- ▶ NEW2: 0.93 IO/SAPS

This means that the IO per second rate was quite similar between model NEW1 and NEW2 which was not influenced by any hardware change.

Effect of increasing throughput

Now we want to compare the measurements NEW1 and NEW2. These environments are based on the same workload profile and running on the same hardware environment (i870 with switchable IASP and SYSBAS). But these environments show different throughputs (605 SAPS for NEW1, 998 SAPS for NEW2) for the same application mix and workload profile. As stated previously in Effect of hardware change the total I/O rate corresponds linear with the throughput.

Now we want to look how the total I/O rate is split up between I/O being made in SYSBAS and I/Os being made in the IASP. Comparing NEW1 and NEW2 you will find the I/O rate of SYSBAS remaining stable at around 184/182 while the I/O rate in the IASP is increasing corresponding with the application workload. The I/O ratio between IASP and SYSBAS is increasing depending on the workload:

- ▶ 2.42 (447/184) for NEW1 (605 SAPS throughput)
- ▶ 4.12 (750/182) for NEW2 (998 SAPS throughput)

K per I/O

The transfer rate per I/O decreases depending on the hardware change (OLD versus NEW1) and the increasing throughput (NEW1 vs. NEW2). There is also a big difference between the transfer rate in SYSBAS comparing with the IASP values. The SYSBAS I/Os are transferring more data (around 13 k per I/O) than the I/O's being made in the IASP (between 6 k and 7 k per IO).

Disk response time, IOP and disk CPU utilization

In our example we provided sufficient resources to cover the I/O requests of the VEKA SAP environment. This is proved by the excellent disk response time and low utilization rate of disk arms and IOP's as shown in Table 3-2.

3.5.3 Disk layout considerations

In order to design a disk layout you must split up the whole data environment into two parts:

- ▶ Data being placed in SYSBAS (temporary storage, IFS files, kernel and work management libraries, SQL package libraries)

- ▶ Data being placed in the IASP (SAP database, sapmnt-directory, Journal receivers)

Note that you must provide sufficient free disk space both in IASP and SYSBAS.

Disk capacity considerations

This section will describe how much disk capacity you should provide for an SAP system running in an IASP. This is done both for existing SAP systems which are running on a non-IASP iSeries landscape and for new SAP systems which need to be sized for an IASP implementation.

New SAP environments

For new SAP environments:

- ▶ In case of IASP size use the SAP quicksizer tool to evaluate the required disk space. Look for a reasonable disk usage (in case of IASP we recommend 60-70%).
- ▶ For sizing the data being placed in SYSBAS refer to Table 3-3, which gives you guidelines about how much temporary storage may be used for specific environments.

Existing SAP environments

If you move an SAP system from a system ASP to a IASP environment without changing the current workload profile you can derive the required disk storage for SYSBAS from the WRKSYSSTS display as shown for customer VEKA in Figure 3-14:

- ▶ Look for 'current unprotected used' while SAP is up and running. This shows how much temporary storage will be requested by the SAP application environment (in our example it was 33 GB).
- ▶ Look for the size of the SAP kernel and work management libraries, SQL packages libraries and IFS directory entries (apart of sapmnt directory, which is placed in the IASP). sum it up and provide enough free space (recommendation: around 50%) to remain flexible in case your workload is growing or your need to place additional data in SYSBAS (for example SAP trace data). Note: An additional SAP instance will also increase the amount of temporary storage.

You also can monitor the size of the SAP database (for example DSPLIB R3<SID>DATA *PRINT) and the size of the sapmnt directory which gives you the total net size being needed for the IASP. Look for a reasonable usage (recommendation: 60-70%).

```

Work with System Status                                PRODSEC1
                                                    11.12.03 16:35:14
% CPU used . . . . . :      33,4  Auxiliary storage:
% DB capability . . . . . :      4,7  System ASP . . . . . :    120,6 G
Elapsed time . . . . . :    00:03:49  % system ASP used . . :    52,4006
Jobs in system . . . . . :     9556  Total . . . . . :      481,4 G
% perm addresses . . . . . :     0,010  Current unprotect used :   33645 M
% temp addresses . . . . . :     0,013  Maximum unprotect . . :   38071 M

Type changes (if allowed), press Enter.

System  Pool  Reserved  Max  -----DB-----  ---Non-DB---
Pool   Size (M)  Size (M)  Active  Fault  Pages  Fault  Pages
  1     914,95   650,70  +++++  0,0   0,0    8,2   9,1
  2    14486,00    0,36    300   31,7  1159  325,6  792,9
  3     819,19    0,00    320   0,0   0,0    0,7   1,2
  4     163,83    0,00     6    0,0   0,0    0,0   0,0

                                                    Bottom

Command
===>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F10=Restart
F11=Display transition data  F12=Cancel  F24=More keys

```

Figure 3-14 WRKSYSSTS - SAP up and running

Disk I/O performance considerations

This gives you a minimum disk requirement from a performance point-of-view, because you have not only to cover the capacity requirements, but also have to provide sufficient disk arms to cover to I/O requests. A lack of disk arms both in IASP and SYSBAS can bottleneck the overall system performance. To avoid such a bottleneck, a minimum number of disk devices is needed for optimal performance depending on the iSeries processor technology and the SAP workload running on that.

New SAP environments

The International SAP/IBM Competence Center (ISICC) in Walldorf published guidelines and recommendations about disk and disk access arm sizing depending on the SAP workload and the iSeries processor technology. These recommendations are referred to an SAP system running exclusively in the system ASP. For details refer to IBM SAP Architecture and Sizing Guidelines, published by the ISICC.

Existing SAP environments

For customer already running the same workload in a non IASP environment (like our customer VEKA), you can take an OS/400 performance measurement and extrapolate the required disk environment for the new system based on a non IASP environment again.

General guidelines for disk layout

As a basis for your calculation take the results both from a capacity and a performance point-of-view. Now you have to split up these disks into one part being configured as a SYSBAS environment and the other one being configured into the IASP. Table 3-3 give you rough guidelines how these configuration can look like. These guidelines were not part of the study done with VEKA and SAP. For new environments you also will find in this table guidelines how much temporary storage an SAP system will need.

Table 3-3 SAP environments - guidelines for disk layout

SAP throughput (User)	Tempspace (SYSBAS)	Disk arm ratio (SYSBAS/IASP)
100	30	1:2
300	50	1:3
500	80	1:4
>1000	100	1:5

Note: these are very rough guidelines, only valid for SAP environments and may vary depending on the SAP release level (our measurement was done with SAP Release 4.6C). Each environment needs to be optimized individually after installation.

3.5.4 Transaction profile and total response time

On Table 3-4 you will find the workload and the response time on the different configurations for the SAP application.

Table 3-4 Customer VEKA - response time overview

	i830 (old) system-ASP	i870 (new) IASP
interactive transactions/day	183,501	202,375
Response time	0.725	0.462
CPU/transaction	266 ms	142 ms
DB request time	285 ms	209 ms



SYSBAS vs IASP configurations

This section details the performance findings of the different applications/scripts that were run with a varying number of disk arms in both the SYSBAS and IASP configurations. In these test runs we were trying to address the answer to the following questions:

- ▶ How big does SYSBAS have to be relative to the IASP(s)? What is the minimal size of SYSBAS?
- ▶ How many DASD arms should I plan for in my IASP?
- ▶ How many DASD arms should I plan for in SYSBAS?

In this chapter we provide the results of the each test pass in summary data and graph formats.

4.1 Why SYSBAS vs IASP configurations are important

As with any system disk configuration the number of disk arms available to the application can have a significant impact to its performance. Too few disk arms will result in fewer disk arms doing more work which in turn can result in longer disk waits and ultimately longer response times to the application.

Results have shown that this is particularly important when it comes to temporary storage in a configured system with IASP. All temporary storage is written to the SYSBAS disk pool. This means that if your application uses a lot of temporary storage you will potentially need more disk arms in the SYSBAS disk pool. Examples might be if you are using SQL applications that need to create temporary indexes/joins or if you are creating large database objects in QTEMP. Conversely, if your application does not use a lot of temporary storage then you can get by with fewer disk arms in the SYSBAS storage pool. Don't forget that the operating system and basic functions occur in the SYSBAS storage pool.

The RPG performance test script results that follow will demonstrate very little usage in the SYSBAS storage pool whereas the SQL performance test scripts purposely utilized the creation of temporary indexes in the SYSBAS storage pool.

4.2 Different DASD configuration scenarios

The DASD configurations used for the two (2) Lab Testing environments(ITSO and Cluster) are discussed here.

4.2.1 ITSO lab environment - iSeries Model 810

The different DASD configuration combinations that were used for the SYSBAS storage pool and the IASP storage pool in the ITSO Lab can be seen in the chart in Figure 4-1. The ITSO Lab system that was utilized had a maximum capacity of 8 DASD arms.

ITSO Lab - iSeries Model 810		SYSBAS Number of DASD Arms				
ITSO Lab - Tested DASD Configurations		0	2	4	6	8
IASP Number of DASD Arms	0					✓
	2		✓	✓	✓	
	4		✓	✓	✓	
	6		✓	✓	✓	

Figure 4-1 ITSO Lab - Model 810 DASD configurations used

4.2.2 Cluster lab environment - iSeries Model 825

The different DASD configuration combinations that were used for the SYSBAS storage pool and the IASP storage pool in the Cluster Lab can be seen in the chart in Figure 4-2. The Cluster Lab system that was utilized had a maximum of 24 DASD arms.

Cluster Lab - iSeries Model 825		Cluster Lab - Tested DASD Configurations								
		SYSBAS Number of DASD Arms								
		0	2	4	6	8	10	12	14	
IASP Number of DASD Arms	0					✓	✓		✓	
	2		✓	✓	✓	✓		✓		
	4		✓	✓	✓	✓				
	6		✓	✓	✓	✓		✓		
	8		✓	✓	✓	✓				
	10									
	12				✓			✓		

Figure 4-2 Cluster Lab - Model 825 DASD configurations used

4.3 Results

This chapter will report the results of the different DASD configuration test runs as they relate to the amount of time elapsed and the average percent busy for each DASD storage pool (SYSBAS and IASP) by application and for each of the two (2) Lab system configurations. Additionally, the amount of time for a failover for each combination of DASD configuration will be reported.

Even though we collected failover data in the sections ITSO Lab - iSeries Model 810 in a testing application and Cluster Lab - iSeries Model 825 in a testing application we didn't chart them. We saw little difference in these failover times. We believe that was due to the fact we didn't have much to recover. We started the multi threaded SQL environment and then waited two minutes and IPL'ed (automated into a CL program). However the queries didn't do enough to invoke SMAPP, etc. So very little variation was seen, unlike what we would expect on a live system.

4.3.1 ITSO Lab - iSeries Model 810 in a testing application

The ITSO Lab system that was utilized had a maximum capacity of 6 DASD arms. when configured with an IASP. The testing applications used in this environment are described in "SQL-based benchmark testing application" on page 17 and "RPG-based benchmark testing application" on page 17.

Summary data

Figure 4-3 below represents the timings and averages for each run.

SYSTEM01 iSeries 9406-810-7430-2469-7430 w/2 Processors, MCHPOOL=0.5GB, *BASE=5.0GB												
IASP # DASD Arms	RPG Secs Plain	RPG Secs w/CMT	RPG Secs w/JRN	DASD % Busy SYSBAS	DASD % Busy IASP	SQL Secs Single	DASD % Busy SYSBAS	DASD % Busy IASP	SQL Secs Multi	DASD % Busy SYSBAS	DASD % Busy IASP	Failover
SYSBAS = 2 DASD Arms												
2	87	236	292	0.00	35.50	939	34.00	14.50	1554	21.50	43.00	414
4	86	235	286	0.00	18.25	981	37.50	9.25	1203	30.00	27.75	404
6	86	234	284	0.00	12.00	950	33.50	7.50	1093	30.00	11.83	406
SYSBAS = 4 DASD Arms												
2	86	237	292	0.00	35.00	683	14.00	21.00	1289	7.75	50.50	431
4	86	234	282	0.00	18.50	700	15.00	12.00	888	11.00	21.25	411
6	86	234	283	0.00	11.83	743	21.00	9.67	821	12.75	16.00	403
SYSBAS = 6 DASD Arms												
2	86	237	291	0.00	36.00	642	7.67	23.50	1259	4.00	54.50	435
4	86	232	283	0.00	17.00	633	7.67	14.50	1053	4.50	31.00	411
6	86	233	283	0.00	12.17	631	7.83	11.50	767	6.67	18.33	413
SYSBAS = 8 DASD Arms												
0	65	168	198	7.63		603	11.00		683	15.38		-

Figure 4-3 Summary Data for the Banking Environment - ITSO Lab

The chart headings/columns is further described as follows:

- ▶ IASP # DASD Arms - This column will list the number of DASD arms allocated to the configured Independent ASP.
- ▶ RPG Secs Plain - This column will contain the elapsed number of seconds the vanilla or plain RPG application took to complete.
- ▶ RPG Secs w/CMT - This column will contain the elapsed number of seconds the RPG application took to complete when journaling and commitment control are active.
- ▶ RPG Secs w/JRN - This column will contain the elapsed number of seconds the RPG application took to complete when journaling is active. This column will contain the elapsed number of seconds the RPG application took to complete when commitment control and journaling are active.
- ▶ DASD % Busy SYSBAS - This column will contain the average percentage busy of the allocated DASD in the SYSBAS storage pool during the RPG Secs w/JRN application test run.
- ▶ DASD % Busy IASP - This column will contain the average percentage busy of the allocated DASD in the Independent ASP storage pool during the RPG Secs w/JRN application test run.
- ▶ SQL Secs Single - This column will contain the elapsed number of seconds the single threaded SQL jobs application took to complete.
- ▶ DASD % Busy SYSBAS - This column will contain the average percentage busy of the allocated DASD in the SYSBAS storage pool during the SQL Secs Single application test run.
- ▶ DASD % Busy IASP - This column will contain the average percentage busy of the allocated DASD in the Independent ASP storage pool during the SQL Secs Single application test run.

- SQL Secs Multi - This column will contain the elapsed number of seconds the multiple threaded (concurrent) SQL jobs application took to complete.
- DASD % Busy SYSBAS - This column will contain the average percentage busy of the allocated DASD in the SYSBAS storage pool during the SQL Secs Multi application test run.
- DASD % Busy IASP - This column will contain the average percentage busy of the allocated DASD in the Independent ASP storage pool during the SQL Secs Multi application test run.
- Secs Failover - This column will contain the elapsed number of seconds it took to fail over the represented DASD configuration from one system to the other system.

Figure 4-4 shows the RPG testing script on the 810 system.

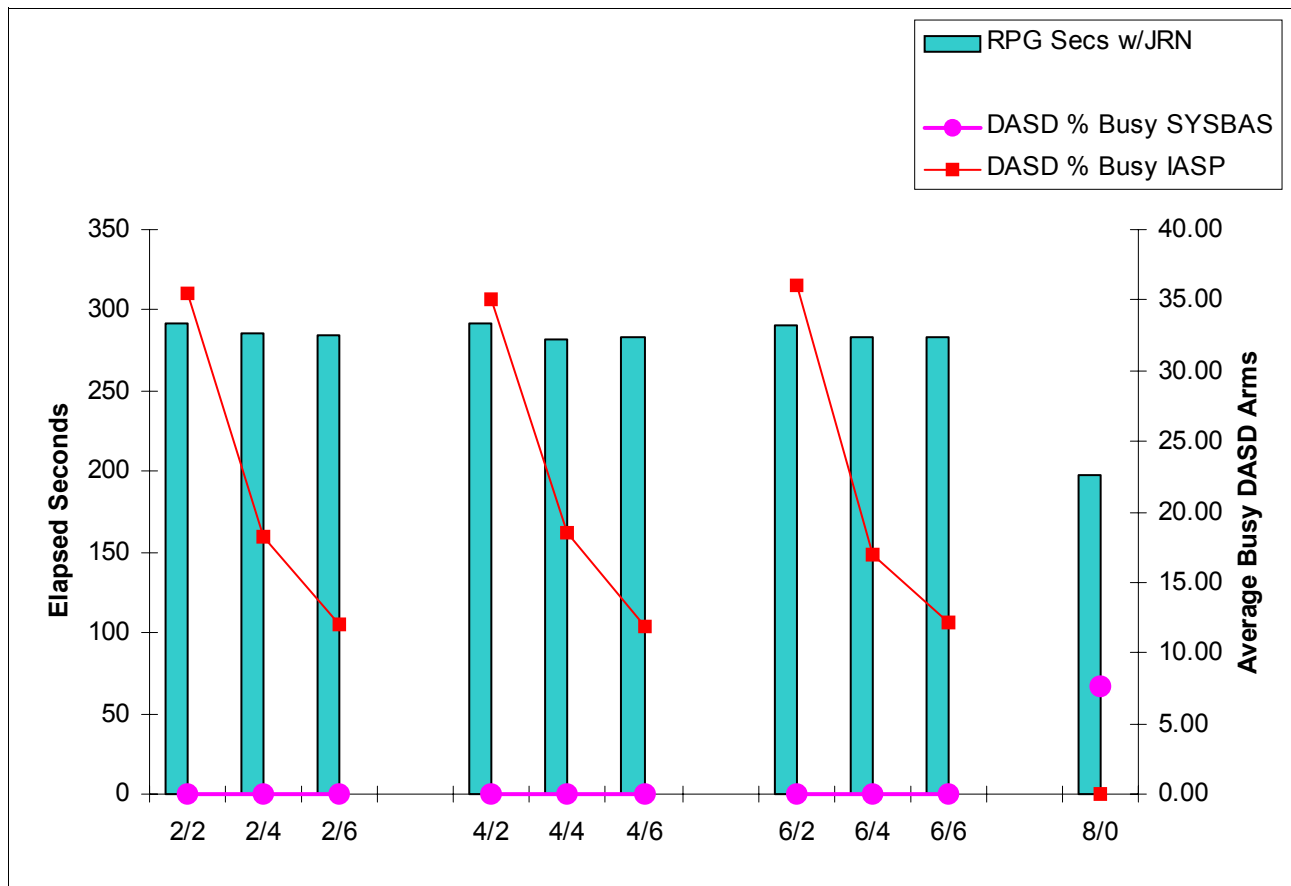


Figure 4-4 RPG with Journaling, number of SYSBAS DASD arms are constant - ITSO Lab (SYSBAS/IASP)

The graph bars show the elapsed time in seconds for completion and the graph lines on the chart show the average DASD percentage busy. From the chart we can see:

- ▶ We can see that the RPG test script is fairly constant in its run completion times across all DASD configurations.
- ▶ We can also see that in a non-IASP environment it was a little quicker in the elapsed time to completion. The DASD percentage busy was lower in the non-IASP environment.
- ▶ We can see that the average DASD percent busy in the IASP was busiest when it only had two (2) DASD units to work with. The significant improvement came when we had a total of four (4) DASD units.

- SYSBAS showed no notable activity on average. This is attributable to the lack of temporary storage utilization, as previously discussed.

Figure 4-5 shows a single threaded SQL series of jobs on the 810 system.

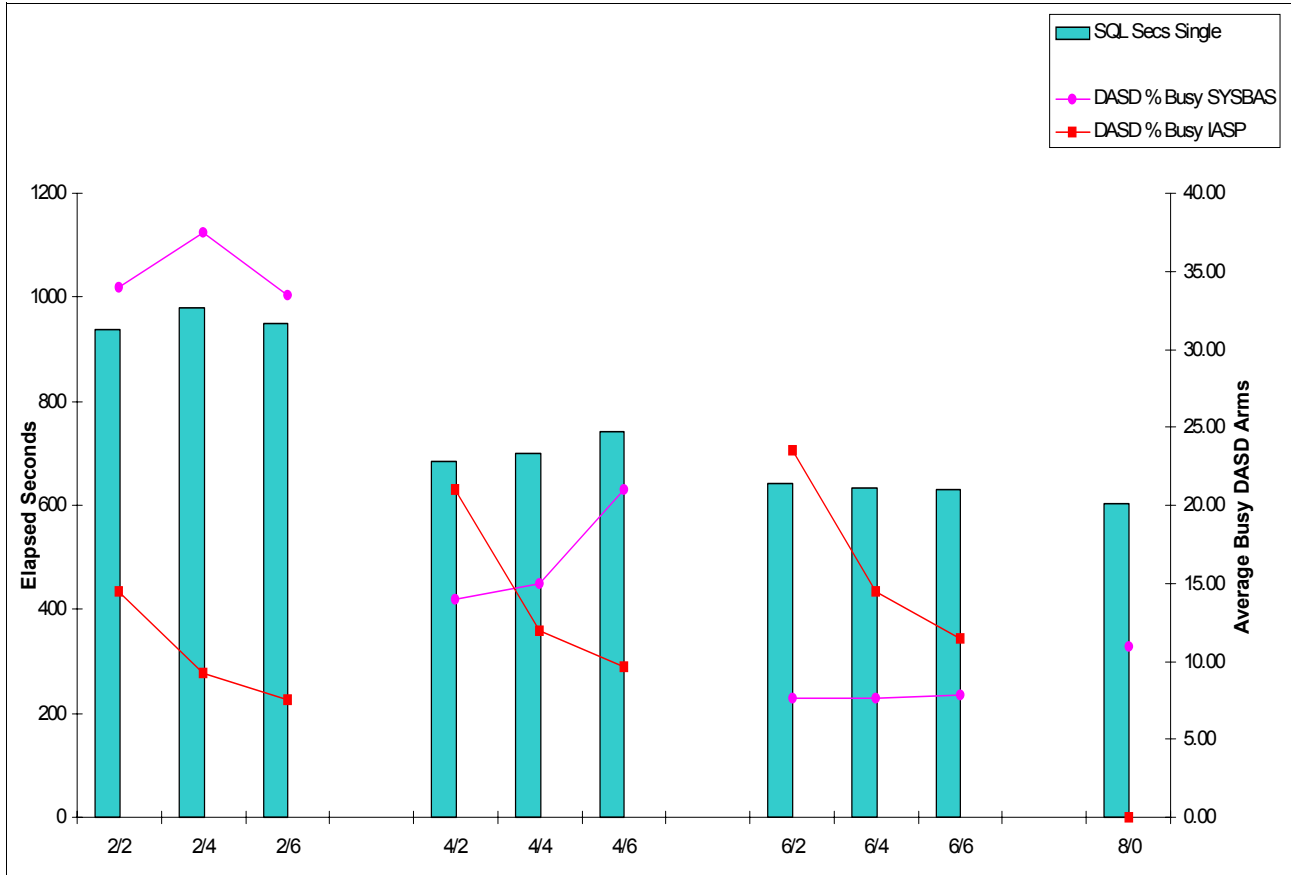


Figure 4-5 Single Threaded SQL, number of SYSBAS DASD arms are constant - ITSO Lab (SYSBAS/IASP)

The graph bars show the elapsed completion time in seconds and the graph lines on the chart show the average DASD percentage busy for each disk pool. We can see that SYSBAS DASD was busiest when it only had two (2) DASD units and it also had the longest elapsed runtime to completion. The average IASP percent busy was not too bad in any tested scenario when we only had two (2) drives in the SYSBAS DASD pool. But, in that case we obviously were spending most of our time waiting on SYSBAS DASD, and were I/O bound there.

Figure 4-6 shows the multiple threaded SQL jobs on the 810 system.




Figure 4-6 Multiple Threaded SQL, number of SYSBAS DASD arms are constant - ITSO Lab (SYSBAS/IASP)

The graph bars show the time in elapsed seconds for completion and the graph lines on the chart show the average DASD percentage busy for each disk pool. We can see that IASP DASD was busiest when it had two (2) DASD units. SYSBAS DASD didn't have the same variance. It seems that the major factor in the multiple threaded SQL job stream is the number of DASD units in the IASP, whereas in the single threaded SQL job stream example, the major bottleneck seemed to be drives in SYSBAS.

4.3.2 Cluster Lab - iSeries Model 825 in a testing application

The Cluster Lab system that was utilized had a maximum capacity of 24 DASD arms. Not all variations were run due to timing limitations. The testing applications used in this environment are described in "SQL-based benchmark testing application" on page 17 and "RPG-based benchmark testing application" on page 17.

Summary data

Figure 4-7 represents the timings and averages for each run.

SYSTEM04 iSeries 9406-825-7418-2473-7418 w/6 Processors, MCHPOOL=1.7GB, *BASE=20.0GB												
IASP # DASD Arms	RPG Secs Plain	RPG Secs w/CMT	RPG Secs w/JRN	DASD % Busy SYSBAS	DASD % Busy iASP	SQL Secs Single	DASD % Busy SYSBAS	DASD % Busy iASP	SQL Secs Multi	DASD % Busy SYSBAS	DASD % Busy iASP	Failover Seconds
SYSBAS = 2 DASD Arms												
2	55	84	123	0.50	24.00	730	34.50	9.00	1000	26.50	36.50	296
4	54	86	113	0.50	10.50	714	37.00	3.50	830	34.00	11.75	476
6	50	91	115	0.00	7.33	716	36.00	2.33	736	36.00	6.00	311
8	50	86	110	1.50	4.75	698	36.50	1.75	712	41.50	4.00	336
SYSBAS = 4 DASD Arms												
2	55	87	125	0.00	21.00	487	12.50	14.00	754	9.50	49.50	311
4	51	84	121	0.00	11.00	460	14.00	5.25	552	17.00	16.25	313
6	50	91	115	1.00	6.67	483	15.50	3.83	619	18.75	7.83	332
8	51	85	109	2.50	4.63	476	16.25	2.88	475	18.50	7.13	311
SYSBAS = 6 DASD Arms												
2	55	90	126	1.00	25.50	478	8.33	13.00	752	7.00	52.00	313
4	51	97	119	1.67	11.50	448	8.67	6.00	552	10.00	18.50	328
6	49	90	115	0.00	7.00	450	8.50	4.33	476	9.00	10.00	307
8	50	87	112	4.33	5.50	434	8.83	2.63	461	13.67	6.50	335
12	50	90	112	3.83	3.58	438	8.67	1.83	455	9.17	4.83	-
SYSBAS = 8 DASD Arms												
0	46	86	110	5.63		444	10.25		456	16.25		-
2	55	84	124	0.88	22.50	470	5.63	13.50	690	6.00	47.00	310
4	52	84	110	1.63	10.00	442	6.25	6.00	539	5.88	17.25	307
6	51	91	115	0.00	13.50	448	6.13	4.67	471	6.50	9.83	306
8	50	86	108	1.88	5.00	436	6.00	3.63	452	6.38	7.75	333
SYSBAS = 10 DASD Arms												
0	47	86	110	4.20		434	8.20		452	12.10		-
SYSBAS = 12 DASD Arms												
2	55	84	123	1.17	22.50	468	3.67	13.50	764	2.42	51.50	-
6	51	93	115	0.00	7.50	442	4.00	4.00	478	5.42	10.50	-
12	50	90	112	0.92	3.08	439	4.17	2.58	452	4.25	5.17	-
SYSBAS = 14 DASD Arms												
0	46	86	111	2.86		433	5.34		443	8.64		-

Figure 4-7 Summary data for the Banking Environment - Cluster Lab

The chart headings/columns is further described as follows:

- ▶ IASP # DASD Arms - This column will list the number of DASD arms allocated to the configured Independent ASP.
- ▶ RPG Secs Plain - This column will contain the elapsed number of seconds the vanilla or plain RPG application took to complete.
- ▶ RPG Secs w/CMT - This column will contain the elapsed number of seconds the RPG application took to complete when commitment control is active.
- ▶ RPG Secs w/JRN - This column will contain the elapsed number of seconds the RPG application took to complete when commitment control and journaling are active.
- ▶ DASD % Busy SYSBAS - This column will contain the average percentage busy of the allocated DASD in the SYSBAS storage pool during the RPG Secs w/JRN application test run.
- ▶ DASD % Busy IASP - This column will contain the average percentage busy of the allocated DASD in the Independent ASP storage pool during the RPG Secs w/JRN application test run.
- ▶ SQL Secs Single - This column will contain the elapsed number of seconds the single threaded SQL jobs application took to complete.
- ▶ DASD % Busy SYSBAS - This column will contain the average percentage busy of the allocated DASD in the SYSBAS storage pool during the SQL Secs Single application test run.
- ▶ DASD % Busy IASP - This column will contain the average percentage busy of the allocated DASD in the Independent ASP storage pool during the SQL Secs Single application test run.
- ▶ SQL Secs Multi - This column will contain the elapsed number of seconds the multiple threaded (concurrent) SQL jobs application took to complete.
- ▶ DASD % Busy SYSBAS - This column will contain the average percentage busy of the allocated DASD in the SYSBAS storage pool during the SQL Secs Multi application test run.
- ▶ DASD % Busy IASP - This column will contain the average percentage busy of the allocated DASD in the Independent ASP storage pool during the SQL Secs Multi application test run.
- ▶ Secs Failover - This column will contain the elapsed number of seconds it took to fail over the represented DASD configuration from one system to the other system.




Figure 4-8 RPG, number of SYSBAS DASD arms are constant - Cluster Lab

Figure 4-8 shows the RPG testing script on the 825 system. The graph bars show the elapsed time in seconds for completion and the graph lines on the chart show the average DASD percentage busy. Unlike the 810, the 825 does show some disk usage in SYSBAS drive pool. The elapsed times for completion were consistent in relation to the number of disk units in the IASP. Additionally, as drives were added to the IASP the average DASD percentage busy was reduced.




Figure 4-9 Single Threaded SQL, number of SYSBAS DASD arms are constant - Cluster Lab (SYSBAS/IASP)

Figure 4-9 shows a single threaded SQL series of jobs on the 825 system. The graph bars show the elapsed completion time in seconds and the graph lines on the chart show the average DASD percentage busy for each disk pool. Once the number of disk arms was increased to at least four (4) disk arms there was consistency with the elapsed completion time in seconds.




Figure 4-10 Multiple Threaded SQL, number of SYSBAS DASD arms are constant - Cluster Lab (SYSBAS/IASP)

Figure 4-10 shows multiple threaded SQL jobs on the 825 system. The graph bars show the elapsed time in seconds for completion and the graph lines on the chart show the average DASD percentage busy for each disk pool. This graph demonstrates the bottleneck on the disk drive usage in SYSBAS when there are only two disk arms. Once this was increased the elapsed time to completion was greatly improved. Thus the number of disk arms that SYSBAS is requiring for the temporary storage is a factor for consideration.

4.3.3 Another view with the total number of disk arms constant

In this section we will examine the data and differences from the perspective of the total number of disk arms being the same but in the different configurations possible between the SYSBAS disk pool and the IASP. First with a total of eight (8) disk arms and then with a total of ten (10) disk arms.

Note: In the following graphs the number of drives in SYSBAS is referenced before the number of drives in the IASP (2-6, 4-4, 6-2, 8-0).




Figure 4-11 RPG and Journaling with a total of 8 DASD arms - ITSO Lab (SYSBAS/IASP)

Figure 4-11 is for the RPG testing script on the 810 looking at it with a total number of eight (8) disk arms.




Figure 4-12 RPG and Journaling with a total of 8 DASD arms - Cluster Lab(SYSBAS/IASP)

Figure 4-12 is for the RPG based benchmark testing application on the 825 looking at it with a total number of eight DASD arms.

Comparing Figure 4-11 and Figure 4-13 we can see in both, that a non-IASP environment runs the quickest in query time. On the 810 there is not a lot of time difference in the queries between the drive configurations, however, the IASP goes down in average DASD arms busy the most from the 6-2, to the 4-4 configuration. We can see that the two best configurations appear to be 2-6 and 4-4. Those would support the goal of a 1-3 configuration or a 1-1 configuration. But since RPG uses no noticeable SYSBAS on this system, we can only see that it runs the best with the most DASD arms in the IASP. The 825 shows the best runtime as a 2-8 configuration showing a 1-4 type configuration. All we can say from this is that RPG shows very little need for arms in SYSBAS.

As we can see all drives SYSBAS performs the best in time for RPG. Looking at DASD Busy, SYSBAS was never enough to register. The IASP had the most improvement when we went from two to four drives in the IASP. So what rule would we recommend? In an environment with everything in RPG, keep as many drives as possible in the IASP, however if we start using the query optimizer, such as with SQL-RPG or OPNQRYF, we would expect different results.

Figure 4-13 is for the single threaded SQL on the 810 looking at it with a total number of eight DASD arms.




Figure 4-13 Single Threaded SQL with a total of 8 DASD arms - ITSO Lab (SYSBAS/IASP)

Figure 4-14 is for the single threaded SQL on the 825 looking at it with a total number of ten DASD arms.




Figure 4-14 Single Threaded SQL with a total of 10 DASD arms - Cluster Lab (SYSBAS/IASP)

Comparing Figure 4-13 and Figure 4-14 we can see that running outside of an IASP is not the advantage as it was with RPG. In the case of the 810 we see a good elapsed runtime and lower DASD busy times in a 4-4 configuration which would tend to support a 1-1 rule. On the 825 we see the best query runtime in a 6-4 configuration that also has low DASD busy times. This would support a 3-2 rule.

Figure 4-15 is for the multiple threaded SQL on the 810 looking at it with a total number of eight DASD arms.




Figure 4-15 Multiple Threaded SQL with a total of 8 DASD arms - ITSO Lab (SYSBAS/IASP)

Figure 4-16 is for the multiple threaded SQL on the 825 looking at it with a total number of ten DASD arms.




Figure 4-16 Multiple Threaded SQL with a total of 10 DASD arms - Cluster Lab (SYSBAS/IASP)

Comparing Figure 4-15 and Figure 4-16 we can see that without IASP is our best elapsed time configuration. In the case of the 810 we see the best performance in the 4-4 configuration, however with the DASD wait times we see with only two drives in the IASP or SYSBAS there is not much other option. In this case a 1-1 type configuration would be shown. On the 825 we see that the best time is in a 6-4 configuration with the quickest query time and the lowest DASD busy times, showing a 3-2 rule.

So what would we recommend? Probably a 1-1 rule or better for DASD in SYSBAS. As with single threaded SQL, temp storage is important to SQL and SYSBAS is where all temp storage is contained.

4.3.4 ITSO Lab - iSeries Model 810 in a data warehousing banking environment

The ITSO Lab system that was utilized had a maximum capacity of 6 DASD arms when configured with an IASP. Not all variations were run due to timing limitations.

Summary data

Figure 4-17 represents the timings and averages for each run.

iASP # DASD Arms	Query 1	Query 2
SYSBAS = 2 DASD		
2	1330	949
4	1273	903
6	1226	878
SYSBAS = 4 DASD		
2	1309	948
4	1232	894
6	1211	876
SYSBAS = 6 DASD		
2	1298	965
4	1260	899
6	1233	877

Figure 4-17 Summary data for the 810 in a data warehousing banking environment

The chart headings/columns is further described as follows:

- ▶ IASP # DASD Arms - This column will list the number of DASD arms allocated to the configured Independent ASP.
- ▶ Query1- This column will contain the elapsed number of seconds the first SQL query took to run.
- ▶ Query2- This column will contain the elapsed number of seconds the second SQL query took to run.

Figure 4-18 shows the 810 elapsed time in seconds for SQL Query1 and SQL Query2. We can see that most of the influence seems to be on DASD arms in the IASP. SQL Query1 was influenced as well by drives in SYSBAS, were SQL Query2 was not influenced as much.




Figure 4-18 Number of SYSBAS DASD arms are constant - ITSO Lab (SYSBAS/IASP)

Figure 4-19 and Figure 4-21 show Query1 and Query2 (respectively) with the data sorted by time. We have turned the 1-3 rule configuration a different color. Based on the query time only, and not on DASD busy information, we can see that the 1-3 rule is doing well, with a 2-3 type configuration doing the best.




Figure 4-19 Elapsed time for Query1 on the 810, Looking at DASD in order of Query Time (SYSBAS/IASP)

Let us take a look at this same data in a normalized form to appreciate the differences.




Figure 4-20 Normalized chart of Query1 on the 810




Figure 4-21 Query2 on the 810, Looking at DASD in order of Query Time

Let us take a look at a normalized version of this data to appreciate the differences.




Figure 4-22 Normalized data from Query2 on the 810. (SYSBAS/IASP)

4.3.5 Cluster Lab - iSeries Model 825 in a Data warehousing Banking Environment

The Cluster Lab 2 system that was utilized had a maximum capacity of 24 DASD arms when configured with an IASP.

Summary data

Figure 4-23 represents the timings and averages for each run.

iASP # DASD Arms	Query 1	Query 2
SYSBAS = 2 DASD		
2	839	702
4	791	670
6	775	640
8	775	632
SYSBAS = 4 DASD		
2	827	702
4	788	661
6	781	654
8	767	634
SYSBAS = 6 DASD		
2	822	687
4	791	665
6	776	653
8	765	640
12	782	632
SYSBAS = 8 DASD		
0	763	613
2	819	703
4	779	656
6	777	635
8	764	641
SYSBAS = 10 DASD		
0	763	617
SYSBAS = 12 DASD		
2	828	701
6	772	661
12	768	622
SYSBAS = 14 DASD		
0	749	608

Figure 4-23 Summary Data for the 825 in a data warehousing banking environment

The chart headings/columns is further described as follows:

- ▶ IASP # DASD Arms - This column will list the number of DASD arms allocated to the configured Independent ASP.
- ▶ Query1 - This column will contain the elapsed number of seconds the first SQL query took to run.
- ▶ Query2 - This column will contain the elapsed number of seconds the second SQL query took to run.

Figure 4-24 shows the elapsed time in seconds for SQL Query1 and SQL Query2. Since we didn't have time to run all the scenarios without drives in the IASP, we connected that data

point to other data in the SYSBAS series with a dotted one, so trends can more easily be seen. We can see that most of the influence seems to be on DASD arms in the IASP.




Figure 4-24 Number of SYSBAS DASD arms are constant - Cluster Lab (SYSBAS/IASP)

Figure 4-25 and Figure 4-27 show Query1 and Query2 (respectively) with the data sorted by time. The best configuration turned out to be the 1-1 configuration. Note that SYSBAS only configurations (no IASP) seem to be doing the best.




Figure 4-25 Query1 on the 825, Looking as DASD in order of Query Time

Let us see a normalized view of the data to appreciate the differences.




Figure 4-26 Normalized data of Query1 on the 825 (SYSBAS/IASP)




Figure 4-27 Query2 on the 825, Looking as DASD in order of Query Time

The following chart is a normalized version of the previous one for better comparison.




Figure 4-28 Normalized version of Query2 on the 825 (SYSBAS/IASP)

4.4 Observations and comments

An IASP configuration is coupled with application type dependencies. Temporary storage usage is major factor in sizing the number of arms in SYSBAS for performance

Like many good things, there is a trade-off. IASPs do have a little bit of performance impact. As you can see from our results, it takes a little longer to run an application in an IASP than it does to run the application in SYSBAS. However, by running in an IASP you get many advantages such as a High Availability solution and the ability to have multiple databases. We feel these advantages easily outweigh the small performance difference. We have also been told that development is working to help minimize this difference in the future as well.

Looking at rules, such as a 1-3 rule (one DASD arm in SYSBAS for every three in an IASP) seems conservative when in environments that use a lot of temporary storage, such as SQL temporary index builds produce. However, remember that we were testing with a small number of overall drives. Where there is room for eight or more drives in SYSBAS and we are still around a 1-3 rule, the results may be very different than what we tested. For what we did test RPG environments not using temp storage will work very well with the 1-3 rule. With SQL environments closer to a 2-3 configuration or even 1-1 configuration is a better way to go.

Also remember that we were testing out one type of application at a time. In a real world environment it would not be uncommon to see a mixture of these types of applications causing different results. Hopefully this data can help give a starting point, but some trial and error may be determined to discover the correct mix of DASD to give the best performance for your application mix.

Also keep in mind that due to hardware limitations on disk drives this study was oriented more towards the iSeries entry market.



User profile considerations

This chapter describes the performance implications of user identification numbers (UID) and group identification numbers (GID) mismatches on the user profiles when we are doing a switchover.

In a clustered environment, a user profile is considered to be the same across servers if the profile names are the same. When switching an IASP between systems as part of a CRG, if the UID and GID are not synchronized for each profile, the UID and GID of each object must be updated on each switch, for each object owned by the mismatched profile. In this chapter we will describe the testings and findings regarding the synchronization of User Profiles.

5.1 Why are user profiles important to IASP Performance

In a clustered environment, a user profile is considered to be the same across servers if the profile names are the same. The name is the unique identifier in the cluster. However, a user profile also contains a user identification number (UID) and a group identification number (GID). To reduce the amount of internal processing that occurs during a switchover, where the independent disk pool is made unavailable (varied off) on one server and then made available (varied on) on a different server, the UID and GID values should be synchronized across the cluster.

When switching an IASP between systems as part of a CRG, if the UID and GID are not in sync for each profile, the UID and GID of each object must be updated on each switch, for each object owned by the mismatched profile. Since this can take a lot of time, it is good to have the UID and GID sync'd for each profile on the systems that use objects in the IASP.

Where is the UID specified? See Example 5-1, the second and third to last options are the UID and GID. The UID and GID numbers are another way of identifying a user or group to a program. For example, the UID and GID numbers are used by programming interfaces in the Integrated File Systems (IFS) environment.

Example 5-1 DSPUSRPRF

```
Display User Profile - *BASIC                                     Page    1
5722SS1 V5R2M0 020719
RCHAS02B 01/07/04 15:28:57
User Profile . . . . . : ZZ0BJ1
Previous sign-on . . . . . :
Sign-on attempts not valid . . . . . : 0
Status . . . . . : *ENABLED
Date password last changed . . . . . : 12/04/03
Password expiration interval . . . . . : *SYSVAL
  Date password expires . . . . . : 06/07/04
Set password to expired . . . . . : *NO
User class . . . . . : *USER
Special authority . . . . . : *NONE
Group profile . . . . . : *NONE
Owner . . . . . : *USRPRF
Group authority . . . . . : *NONE
Group authority type . . . . . : *PRIVATE
Supplemental groups . . . . . : *NONE
Assistance level . . . . . : *SYSVAL
Current library . . . . . : *CRTDFT
Initial program . . . . . : *NONE
  Library . . . . . :
Initial menu . . . . . : MAIN
  Library . . . . . : *LIBL
Limit capabilities . . . . . : *NO
Text . . . . . :
Display sign-on information . . . . . : *SYSVAL
Limit device sessions . . . . . : *SYSVAL
Keyboard buffering . . . . . : *SYSVAL
Storage information:
  Maximum storage allowed . . . . . : *NOMAX
  Storage used . . . . . : 12
  Storage used on independent ASP . . . . . : *NO
Highest scheduling priority . . . . . : 3
Job description . . . . . : QDFTJOB
  Library . . . . . : QGPL
Accounting code . . . . . :
```

```

Message queue . . . . . : ZZOBJ1
  Library . . . . . : QUSRSYS
Message queue delivery . . . . . : *NOTIFY
Message queue severity . . . . . : 00
Output queue . . . . . : *WRKSTN
  Library . . . . . :
Printer device . . . . . : *WRKSTN
Special environment . . . . . : *SYSVAL
Attention program . . . . . : *SYSVAL
  Library . . . . . :
Sort sequence . . . . . : *SYSVAL
  Library . . . . . :
Language identifier . . . . . : *SYSVAL
Country or region identifier . . . . . : *SYSVAL
Coded character set identifier . . . . . : *SYSVAL
Character identifier control . . . . . : *SYSVAL
Locale job attributes . . . . . : *SYSVAL

```

Display User Profile - *BASIC

```

Page      2
5722SS1 V5R2M0 020719
RCHAS02B 01/07/04 15:28:57

```

```

User Profile . . . . . : ZZOBJ1
Locale . . . . . : *SYSVAL
User options . . . . . : *NONE
Object auditing value . . . . . : *NONE
Action auditing values . . . . . : *NONE
User ID number . . . . . : 6000001
Group ID number . . . . . : 8000001
Home directory . . . . . : /home/ZZOBJ1

```

* * * * * E N D O F L I S T I N G * * * * *

For library objects that have a profile mismatch for the UID and/or GID, we will get the message CPC2206 logged. See Example 5-2 for a sample message CPC2206 we received for just such a mismatch.

Example 5-2 Message CPC2206

```

CPC2206  Completion      00 12/05/03 10:24:13.548960 QSYCHONR  QSYS      05E3  QSYIASP  QSYS      *STMT
To module . . . . . : QSYXFER
To procedure . . . . . : XferFromDangler_QSYS__FPvT1i
Statement . . . . . : 2282
Thread . . . . . : 00000002
Message . . . . . : Ownership of object ZZ00655 in ZZOBJ1 type *FILE changed.
Cause . . . . . : The ownership of object ZZ00655 in library ZZOBJ1 type
                  *FILE has changed.

```

We received no equivalent message for IFS objects with a profile mismatch for UID and/or GID, however with more mismatches, the vary on did take increasing amounts of time.

5.2 Data used for the User Profile testing

In this section we will look at the raw data gathered during the testing.

Besides measuring the failover and switchover times as discussed in “Failover times” on page 21 and also in “Switchover times” on page 22, we also measured the time that the vary on joblog took on the target system of the Failover or Switchover. The name of the vary on jobs for cluster is QCSTVRYDEV. Because the Vary on joblog is not associated with anything on the system the IASP is moving from, it should always have less time than the failover or

switchover itself. Since the correction to the object occurs in the vary on process, we would expect to see most of the time spent there for the failover or switchover.

Figure 5-1 shows the results we found in seconds for the Failover and Switchover testing going from the 820 to the 810 for library objects. We have done this for both UID mismatches and for UID and GID mismatches. The measurements were converted to seconds for the chart. The numbers on the top of each column shows the number of objects that were mismatched(0,1000,2000,3000,4000,100000).

	0	1000	2000	3000	4000	100000
UID Mismatch Failover	286	421	486	619	644	10686
UID Mismatch Failover Vary on Only	151	314	413	491	548	9938
UID/GID Mismatch Failover	286	446	524	613	647	-
UID/GID Mismatch Vary on Only	151	347	429	485	552	-
UID Mismatch Switchover	264	340	439	494	546	10036
UID Mismatch Switchover Vary on Only	242	320	415	471	525	9938
UID/GID Mismatch Switchover	264	344	436	492	553	-
UID/GID Mismatch Switchover Vary on Only	242	324	424	474	527	-

Figure 5-1 820 -> 810 time in seconds for Failover and Switchover for Library objects

Figure 5-2 shows the results we found in seconds for the Failover and Switchover testing going from the 810 to the 820 for library objects. We have done this for both UID mismatches and for UID and GID mismatches. The measurements were converted to seconds for the chart.

	0	1000	2000	3000	4000	100000
UID Mismatch Failover	263	443	536	599	668	11090
UID Mismatch Failover Vary on Only	46	226	324	386	455	10342
UID/GID Mismatch Failover	263	435	538	603	693	-
UID/GID Mismatch Failover Vary on Only	46	221	322	387	500	-
UID Mismatch Switchover	172	307	372	552	619	10591
UID Mismatch Switchover Vary on Only	49	214	325	414	483	10492
UID/GID Mismatch Switchover	172	430	510	548	636	-
UID/GID Mismatch Switchover Vary on Only	49	212	327	410	487	-

Figure 5-2 810 -> 820 time in seconds for Failover and Switchover for Library objects

Figure 5-3 shows the results we found in seconds for the Failover and Switchover testing going from the 820 to the 810 for IFS objects. We have done this for both UID mismatches ad for UID and GID mismatches. The measurements were converted to seconds for the chart.

	0	1000	2000	3000	4000	100000
UID Mismatch Failover	251	271	303	328	363	1449
UID Mismatch Failover Vary on Only	148	177	210	232	268	1327
UID/GID Mismatch Failover	251	271	305	357	387	1358
UID/GID Mismatch Vary on Only	148	178	209	231	261	1289
UID Mismatch Switchover	191	204	229	258	286	1375
UID Mismatch Switchover Vary on Only	154	180	208	236	269	1279
UID/GID Mismatch Switchover	191	204	229	258	285	1373
UID/GID Mismatch Switchover Vary on Only	154	181	208	234	265	1279

Figure 5-3 820 -> 810 time in seconds for Failover and switchover for IFS objects

Figure 5-4 shows the results we found in seconds for the Failover and Switchover testing going from the 810 to the 820 for IFS objects. We have done this for both UID mismatches ad for UID and GID mismatches. The measurements were converted to seconds for the chart.

	0	1000	2000	3000	4000	100000
UID Mismatch Failover	258	283	350	378	410	1288
UID Mismatch Failover Vary on Only	49	71	94	127	162	1291
UID/GID Mismatch Failover	258	320	349	383	411	1442
UID/GID Mismatch Failover Vary on Only	49	71	100	118	161	1395
UID Mismatch Swichover	178	200	233	259	282	1386
UID Mismatch Swichover Vary on Only	40	65	97	123	144	1293
UID/GID Mismatch Swichover	178	204	231	258	289	1389
UID/GID Mismatch Swichover Vary on Only	40	65	94	123	149	1296

Figure 5-4 810 -> 820 time in seconds for Failover and swichover for IFS objects

The next two charts show a comparison of object processing for profile mismatch when we had two processors defined on the 820 and with three processors defined on the 820.

Figure 5-5 shows Swichover from the 820 to the 810 when IFS objects had a mismatch. This can help illustrate if having an extra processor will help the system that the IASP is switching from.

	0	100000
2 Processor UID Mismatch	217	1380
2 Processor UID Mismatch Vary on Only	123	1283
2 Processor UID/GID Mismatch	217	1376
2 Processor UID/GID Mismatch Vary on Only	123	1284
3 Processor UID Mismatch	191	1375
3 Processor UID Mismatch Vary on Only	154	1279
3 Processor UID/GID Mismatch	191	1373
3 Processor UID/GID Mismatch Vary on Only	154	1279

Figure 5-5 820 -> 810 Swichover for IFS objects with a UID and UID/GID profile mismatch with both two and three processor data for the 820

Figure 5-6 shows Swichover from the 810 to the 820 when IFS objects had a mismatch. This can help illustrate if having an extra processor will help the system that the IASP is switching to.

	0	100000
2 Processor UID Mismatch	223	1377
2 Processor UID Mismatch Vary on Only	142	1286
2 Processor UID/GID Mismatch	223	1388
2 Processor UID/GID Mismatch Vary on Only	142	1298
3 Processor UID Mismatch	178	1386
3 Processor UID Mismatch Vary on Only	40	1293
3 Processor UID/GID Mismatch	178	1389
3 Processor UID/GID Mismatch Vary on Only	40	1296

Figure 5-6 810 -> 820 Swichover for IFS objects with a UID and UID/GID profile mismatch with both two and three processor data for the 820

The results in the next sections that had data collected for 100,000 objects with mismatches will usually be shown in two charts, one with and one without the 100,000 objects referenced. This is because when the 100,000 objects are referenced it is hard to see the characteristics of the first four mapped points.

5.3 Failover vs. switchover

Figure 5-7 and Figure 5-8 will both be looked at together. Both figures are library failover and switchover. Figure 5-7 is specific to the 820 failover and switchover to the 810. On the other hand Figure 5-8 is specific to the 810 failover and switchover to the 820.

In Figure 5-7 and Figure 5-8 we see that as the number of mismatched objects go up, the time also increases. This is probably what most people would indeed guess! We can also see that there is a much greater change in the time between the whole failover vs. the whole switchover, than the failover vary on vs. the switchover vary on.

In both Figure 5-7 and Figure 5-8 we see that there is not much difference between the vary on joblogs. In most cases the Failover Vary on will take a little longer than the Switchover Vary on.





Figure 5-7 820 -> 810 Library Failover and Switchover for UID Mismatch.

On Figure 5-7 the difference between the top chart and the bottom one is that the bottom one includes the 100,000 data point.




Figure 5-8 810 -> 820 Library Failover and Switchover for UID Mismatch

On Figure 5-8 the difference between the top chart and the bottom one is that the bottom one includes the 100,000 data point.

Looking at Figure 5-9 and Figure 5-10 we can see that IFS Failover and Switchover is quite similar to the Library Failover and Switchover. For the most part, there seems to be more variation between the whole Failover time and the whole Switchover time than between the

Failover Vary on and the Switchover vary on. So we can assume that the extra time that Failover takes instead of Switchover in this case is due in the hardware move and not the process of varying the IASP back on.

Now it should be remembered, that this is NOT a typical failover situation. Nothing was using the IASPs so there was little recovery to be done. A typical failover vary on would also have to deal with cleaning up objects that were in use at the time that the system went down. Since we were only testing the UID and GID mismatch issues for these situations, we didn't want to introduce other variables, so the IASPs were not in use.

We see in Figure 5-9 the switchover job took longer than the failover. We believe this is a data anomaly based on our other testing and if we ran this particular scenario more times and averaged out the results we would see that failover generally takes more time than switchover.




Figure 5-9 820 -> 810 IFS Failover and Switchover for UID Mismatch

On Figure 5-9 the difference between the top chart and the bottom one is that the bottom one includes the 100,000 data point.




Figure 5-10 810 -> 820 IFS Failover and Switchover for UID Mismatch

On Figure 5-10 the difference between the top chart and the bottom one is that the bottom one includes the 100,000 data point.

5.4 Comparing UID vs. UID/GID profile mismatches

Figure 5-11 shows us all library vary on Switchover jobs. Both UID and UID/GID mismatches are looked at. We can see looking back at Example 5-2 on page 79 that there is no mention if the UID didn't match, or the GID didn't match. Also, there is only the one message sent even if both the UID and GID are mismatched. The conclusion we can draw from this is that it doesn't matter if you have only the UID or GID mismatched, or if both the UID and GID are mismatched. The correction during the IASP vary on will take the same amount of time. Of course the best option is to ensure that there is no mismatch. See "User profiles" on page 18 for how to check for mismatches, and "OK, how do I fix it" on page 93 for how to correct them.




Figure 5-11 Library objects with a UID vs. UID/GID profile mismatch for the vary on switchover.

In Figure 5-12 we can see much the same story as Figure 5-11. That is, little difference between the objects with a UID mismatch and the objects with a UID and GID profile mismatch. So IFS and Library objects work about the same in the fact that the type of mismatch does not matter. However see "Library vs. IFS" on page 89 for how IFS vs. Libraries do differ.

One thing that we can note is that below 100,000 objects the 820 is always faster than the 810. At the 100,000 object mismatch mark, it seems we pretty much even out. We assume this is because we are more DASD bound, then we are CPU bound.




Figure 5-12 UID vs. UID/GID Profile Mismatch for IFS switchover vary on

5.5 Library vs. IFS

Figure 5-13 is a very interesting one. The part that does not show the 100,000 objects seems fairly normal. We can see that the 820 is faster than the 810 and that library objects take

longer to update the UID on than IFS objects do. We know that as in Example 5-2 on page 79 we get messages for library objects with a mismatch, however there is no equivalent message for IFS. This may be what causes the differences in time.

Looking down at the data that does include the 100,000 data point we can see that the larger the number of objects with a UID or GID profile mismatch, the greater the gap becomes between Library and IFS processing. We can easily conclude that IFS objects are able to be corrected quicker than Library objects.




Figure 5-13 Library vs IFS Mismatches shown through objects with a UID Mismatch in the user profile for a switchover vary on job

5.6 How much difference a processor makes

How much difference does a processor make? Figure 5-14 and Figure 5-15 make it quite clear that a processor makes more difference at the low end. If you have few objects with a UID or GID profile mismatch it will help. However, by the time we get to the 100,000 object range, the task seems to be DASD bound and a processor doesn't do us a lot of extra good.




Figure 5-14 820 -> 810 2 processors vs. 3 processors in the 820, in an IFS Switchover for objects with a UID profile mismatch for the vary on job only




Figure 5-15 810 -> 820 2 Processors vs. 3 Processors in the 820, in an IFS Switchover for objects with a UID profile mismatch for the vary on job only

We can say that since Figure 5-15 shows a greater initial separation than Figure 5-14, that more processors will help a little more on the system getting the IASP, than the one that the IASP is coming from.

5.7 OK, how do I fix it

The purpose of this section is to help you understand that the best scenario is to have a user profile's UID and GID sync'ed up. So once this has been accomplished, the obvious question is, what do we do to go about fixing the situation?

"Checking for profile UID/GID mismatch" on page 19 shows how to find profiles that do have an existing UID, GID, or both mismatch. Once we have the list, we can start on fixing the situation.

We did some testing and found that a CHGUSRPRF will work if:

1. No active thread is using the user profile.
2. The user profile owns no IFS objects.

If the profile is associated with an active thread, Example 5-3 will be shown.

Example 5-3 Message CPF22DE

Additional Message Information

Message ID :	CPF22DE	Severity :	40
Message type :	Escape	Time sent :	14:02:18
Date sent :	01/09/04		

Message : Not allowed to change the UID or GID of user profile
ZZOBJ7.
Cause : A job is active which has ZZOBJ7 as the process profile or
one of the process group profiles. The UID and GID of a profile may not be
changed when the profile is active in a job.
Recovery . . . : Wait until all jobs with the process profile or a process
group profile of ZZOBJ7 complete and try the request again. If you are using
the Change User Profile UID or GID (QSYCHGID) API, then you may need to take
the system to restricted state (end all subsystems) to change the UID or
GID. See the OS/400 APIs topic in the Information Center book,
<http://www.iseries.ibm.com/infocenter>, for details.

Bottom

Press Enter to continue.

F3=Exit F6=Print F9=Display message details
F10=Display messages in job log F12=Cancel F21=Select assistance level

If the profile owns an IFS object, and you try to do a CHGUSRPRF the message shown in
Example 5-4 will be shown.

Example 5-4 Message CPF22DC

Additional Message Information

Message ID : CPF22DC Severity : 40
Message type : Escape
Date sent : 01/09/04 Time sent : 14:03:24

Message : Not allowed to change UID of the user profile.
Cause : The UID of a user profile may not be changed when the
profile is the owner of an object in a directory.
Recovery . . . : Change the value of the UID parameter to *SAME and try the
command again.

Bottom

Press Enter to continue.

F3=Exit F6=Print F9=Display message details
F10=Display messages in job log F12=Cancel F21=Select assistance level

There is no easy way around these messages via 5250. First, a time to do the change would
have to be found when the profile isn't active. Second the IFS objects would have to be
reassigned first. It might be easier to create a temporary profile and assign the IFS objects to
the profile, then after the change is done, move them back.

In our testing, we did find that Management Central was able to change the UID and GID of a
profile when it owned both library and IFS objects. That makes Management Central the easy
way to go for this.

Note that Management Central is the prefer way to fix this situation.

Tip: Management Central can also assist with synching UID's and GID's for user profiles. See the Info Center article at:

<http://publib.boulder.ibm.com/iserics/v5r1/ic2924/info/rzaih/rzaihsynchuidgid.htm>



Part 3

Conclusions and recommendations

In this part we describe the conclusions and recommendations of our findings from the tests that we conducted. It represents a summary of our findings.



Conclusions and recommendations

Throughout the previous sections of this redpaper, we have introduced you to a number of different environments that illustrate the performance implications of configuring Independent ASPs.

We also showed the measured impact of some of these factors in a real customer environment and indicated a few of the reasons for our measured findings.

In this section, we will revisit some of these issues and also share with you a few general recommendations that could further assist you in understanding the performance implications of implementing Independent ASPs.

6.1 General performance findings

Let us revisit the questions that we wanted to answer with this study and let us see what are our conclusions from this study:

- ▶ *Question:* How big does SYSBAS have to be relative to the IASP(s)? What is the minimal size of SYSBAS? How many arms should I plan for in my IASP? How many arms should I plan for in SYSBAS?

An IASP configuration is coupled with application type dependencies. Temporary storage usage is major factor in sizing the number of arms in SYSBAS for performance.

In order to design a disk layout and for your storage requirements you must split up the whole data environment into two parts:

- Data being placed in SYSBAS (temporary storage, IFS files, work management libraries, and SQL package libraries)
- Data being placed in the IASP (Database, Journal receivers)

Note: Provide sufficient free disk space both in IASP and SYSBAS. There should be a minimum number of disk arms for good performance (for both SYSBAS and IASP). This should be based on the normal system sizing tools.

Like many good things, there is a trade-off. IASP's do have a little bit of performance impact. As you can see from our results, it takes a little longer to run an application in an IASP than it does to run the application in SYSBAS. However, by running in an IASP you get many advantages such as a High Availability solution and the ability to have multiple databases. We feel these advantages easily outweigh the small performance difference. As an example of the benefits of implementing applications using Independent ASPs we have illustrated a real customer case with SAP R/3.

Another of our conclusions were: the major factor in the multiple threaded SQL job stream seems to be the number of DASD units in the IASP, whereas in the single threaded SQL job stream, the major bottleneck seemed to be the number of disk drives in SYSBAS.

There are applications that have no correlation to the ratio of storage between SYSBAS and the IASP. Rather it is directly correlated to the number of arms in the IASP.

Looking at the ratio between SYSBAS and IASP there is no apparent conclusion to this ratio. There are applications dependencies, such as the use or not use of a lot of temporary storage that would make this ratio change and difficult to generalize. We suggest that every environment has to be treated separately and if customers use the approach used by VEKA shown in Chapter 3, "SAP - Customer scenario" on page 27 they will be better prepared for this change.

- ▶ *Question:* How does user profile differences affect failover / switchover time? How does the user profile synchronization affects the time?

In a clustered environment, a user profile is considered to be the same across servers if the profile names are the same. When switching an IASP between systems as part of a CRG, if the UID and GID are not synchronized for each profile, the UID and GID of each object must be updated on each switch, for each object owned by the mismatched profile.

Regarding the User Profiles UID and GID mismatches we have concluded:

- The synchronization of UID and GID mismatches of library objects take longer than IFS objects.
- It is not the number of user profiles that matters, it is the number of objects regardless of their size that makes the synchronization longer or not.

Our recommendation is to enforce a strategy in your shop to minimize this mismatches.

- ▶ *Question:* What kind of failover / switchover time can I expect with different configurations?

This was not the main focus of the study but we did find out that disk arm ratios between SYSBAS and IASP had zero effect - (no positive - no negative) - on the failover (switchover) performance.

6.2 Other performance recommendations

This are some useful recommendations to take into consideration:


The time required to varyon an ASP group is affected by the number of files in SYSBAS. Each ASP group represents a separate instance of a database and each contains its own set of cross-reference and catalog files. Since an ASP group database also includes the files in SYSBAS, varyon merges the information from the cross-reference and catalog files in SYSBAS into the cross-reference and catalog files of the ASP group. The more records that are contained in the files in SYSBAS, the longer it takes to perform the database merge function of varyon processing. To achieve the fastest varyon times possible, one should limit the files that exist within SYSBAS. Some other object types can have entries in the cross-reference and catalog files but generally do not have enough entries that their placement in SYSBAS has much affect on varyon processing time.

Creating and deleting files in SYSBAS requires more processing with IASPs. Every change to the cross-reference and catalog files in the system ASP causes a similar operation to be performed for every ASP group that is in an AVAILABLE state. As was described in the previous paragraph, this overhead can be avoided by limiting the files that are in SYSBAS.

Using the Set ASP Group (SETASPGRP) command or job description parameter allows ASP group database operations to be performed within the job that has done the SETASPGRP. Using SQL to connect to an ASP group database results in a local socket connection to a local QRWTSRVR job and the database operations being performed in that job. This can introduce considerable overhead and result in degraded performance. Even if an application contains a CONNECT statement, preceding the CONNECT with a SETASPGRP or having it set through the job description parameter results in the database operations being performed in the job issuing the CONNECT.

Just as SMAPP can be used to limit the time spent in access path recovery during an IPL following a system failure, SMAPP can be used to limit the time spent for the same function during varyon of an IASP. The SMAPP setting is tailorable by ASP group. The ASP groups that are most critical should be given the lower SMAPP settings.

If feasible yet not too inconvenient, separate mission critical data from non-critical data into different ASP groups or UDFS ASPs. If there is a system failure, the IASPs containing mission critical data can varyon first. Those IASPs containing non-critical data can varyon sometime later when they will not slow the varyon of the mission critical IASPs.



Supported and unsupported objects in IASPs

This appendix provides the list of supported and unsupported objects in IASPs.

Supported object types

*ALRTBL	*FIFO	*MODULE	*QRYDFN
*BLKSF	*FILE	*MSGF	
*BNDDIR	*FNTRSC	*MSGQ	*SCHIDX
*CHTFMT	*FNTTBL	*NODGRP	*SPADCT
*CHRSF	*FORMDF	*NODL	*SQLPKG
*CLD	*FTR	*OVL	*SQLUDT
	*GSS	*PAGDFN	*SRVPGM
*CMD	*IGCDCT	*PAGSEG	*STMF
*CRQD	*JOBQ	*PDG	*SVRSTG
*CSI	*JRN	*PGM	*SYMLNK
*DIR	*JRNRCV	*PNLGRP	*TBL
	*LIB	*PSFCFG	*USRIDX
*DTAARA	*LOCALE	*QMFORM	*USRQ
*DTADCT	*MEDDFN	*QMQR	*USRSPC
*DTAQ	*MENU		*VLDL

Figure A-1 Supported object types

Restrictions for supported object types: *FILE - Database files that are either multi-system database files, or that have fields in them that are DataLink fields that are created as Link Control, cannot be located in an independent disk pool

Unsupported object types

*AUTHLR	*CTLD	*IGCTBL	*NTBD
*AUTL	*DDIR	*IPXD	*NWID
*CFGL	*DEVQ	*JOBQ	*NWSQ
*CNL	*DOC	*JOBSCD	*OUTQ
*COSD	*DSTMF	*LIND	*PRDAVL
*CRG	*EDTD	*MODD	*USRPRF
*CSPMAP	*EXITRG	*M36	*SOCKET
*CSPTBL	*FLR	*M36CFG	*SSND
	*IGCSRT		*S36

Figure A-2 Unsupported Object Types

Related publications

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

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 105. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM eServer iSeries Independent ASPs A Guide to Moving Applications to IASPs*, SG24-6802

Other publications

These publications are also relevant as further information sources:

- ▶ *iSeries Backup and Recovery*, SC41-5304

Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ Information Center
<http://publib.boulder.ibm.com/series/v5r2/ic2924/info/rzaly/rzalyapplications.htm>

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Independent ASP Performance Study

on the IBM *e*server iSeries Server



**Understand how IASP
drive configuration
can affect your
application**

**Understand user
profile mismatch
implications**

**Learn about an SAP
true life
implementation**

Independent ASPs were announced and made available in V5R1. The initial introduction of IASPs was limited to IFS objects. As such, iSeries customers started using this concept for Domino, Web serving applications, and Integrated xSeries Adapters servers. In V5R2 some major enhancements were made to the support of IASP. Now they support traditional library and object structure in the QSYS.LIB file system used by the majority of iSeries applications. For this reason many customers have been asking what the performance considerations are when you move your application to IASPs.

The aim of this Redpaper is to provide you with an understanding of the performance implications when you decide to move your applications to Independent ASP. This redpaper should be read in conjunction with the existing IBM Redbook *IBM eServer iSeries Independent ASPs A Guide to Moving Applications to IASPs*, SG24-6802. It is the purpose of this redpaper to complement the existing IBM Redbook using real customer scenarios to illustrate some performance considerations.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:
ibm.com/redbooks**