



LPAR For Decision Makers

*Why Logical Partitioning on IBM @server
pSeries Servers Matters to Your Business*

November 12, 2002

I. Introduction

Logical Partitioning (LPAR) is a system architecture approach that promises to revolutionize corporate IT infrastructures and have a huge impact on business operations and equipment purchase decisions. Originally developed for mainframe computers, LPAR allows the division of a single server into several completely independent “virtual” servers or partitions. LPAR on IBM @server™ pSeries™ POWER4™ -based servers is supported by AIX® 5L™, IBM’s UNIX® operating system, and by the Linux® operating system available from one or more Linux Distributors.

Each partition is independent of the operations occurring within other partitions, can run its own version of the operating system, and have processor, memory, and I/O resources dedicated to its exclusive use. Dynamic LPAR allows these resources to be increased or decreased, with no interruption on operations. IBM has taken the flexibility offered by this technology out of the mainframe computer center and put it into the low cost POWER4-based pSeries servers running AIX 5L Version 5.2.

This paper is a non-technical introduction to the concepts and advantages of LPAR. It discusses the business benefits that might result from the inclusion of LPAR-capable servers in an overall IT infrastructure. It points out particular business scenarios and applications where LPAR might be used to good advantage. Lastly, it compares and contrasts the implementations offered by different vendors. The goal of this paper is to help non-IT managers and executives better understand LPAR and how it can improve business results.

II. Business and IT Benefits of LPAR

The purpose of a company’s IT infrastructure is to allow critical business processes and functions to operate productively and efficiently, with minimum investment and ongoing operational cost. Improvements to an IT infrastructure are always a trade off between how much additional business productivity can be achieved versus the equipment and personnel investment necessary to achieve that productivity increase.

LPAR can contribute significantly to either the reduction of IT infrastructure cost or the improvement in workload capacity in the following ways.

- Reduced systems management requirements and improved system administration productivity
- Reduced use of floor space or rack space
- Reduced power consumption
- Reduced software licensing fees
- Reduced equipment service and maintenance costs
- Improved business operation flexibility
- Improved security

The following scenarios illustrate these advantages:

1. Cost effective handling of peak workloads

Most servers are generally sized to meet the demands of the peak workloads they are expected to run, even though the average use of those systems may be far less. This results in systems being underutilized most of the time, yet strained to capacity at those times when the workload peaks occur. If the workload peaks are known and can be scheduled, then a system administrator can add workloads or applications to take advantage of the unused server capacity. If the workload peaks are not regular or predictable, workload scheduling can become a problem that impacts business operations.

Dynamic LPAR-capable servers allow resources to be quickly and easily shifted to the workloads that require them. This makes meeting the peak workload demands much easier.

With dynamic LPAR, the server can be subdivided into partitions each with their own copy of an operating system and key workloads assigned to different partitions. Processor and memory resources can be dynamically allocated to each partition as the applications in that partition demand, and then released for use elsewhere by applications in other partitions. And when the peak load for which the server was sized occurs, the partitions can be reduced to contain a minimum of resources or even closed in order to release the resources needed. With dynamic LPAR-capable servers, a partition reboot is not required when resources are reallocated, and there is no impact on system availability.

The net result is improved computing asset utilization which translates into either the need for a smaller sized system with a smaller price tag, or the ability to efficiently shift more workloads to that server and perhaps delay or even eliminate the need for additional servers.

There is an additional benefit to logical partitioning. Since the partitions are completely independent, a failure of a component assigned to that partition doesn't affect the operations of other partitions. A software fault which brings its partition down has no effect on the applications operating in different partitions. You get much of the reliability of independent physical servers with the flexibility benefits of being able to shift physical resources to the points where they are most needed.

2. Mixed testing and production environments

One of the difficulties in introducing new applications, upgrading existing applications, or moving to new versions of operating systems is the potential for disruption of the business processes. To prevent this disruption, many people insist on long certification cycles before modifications are migrated to the existing business process environment. However, these implementation delays can also extract their own penalties.

Partitioning provides a means for setting aside a portion of the system resources to be used for testing new versions of applications and operating systems while the production environment continues to run. This avoids the expense of buying additional servers for test purposes only, and also gives higher confidence that the test versions will migrate smoothly into production because they are tested on the actual production hardware

system. This ability to reduce the risk of adversely impacting business operations when bringing new application or operating system upgrades or fixes on-line is a key benefit of LPAR.

3. Server consolidation requiring data or application independence

One of the risks of running multiple workloads on a multitasking server is the potential of a failure on one application affecting the operation, security, or data integrity of another application. For this reason, many customers outsourcing IT functions to Application Service Providers insist on complete isolation of their workload from anything else. This, of course, increases costs because when separate servers are used for each individual application, unused excess capacity can result. The individual systems take up more physical space, consume more power and increase the air conditioning demand, require more personnel and system management, increase maintenance costs, must be individually backed up, and are typically underutilized.

Running applications in separate partitions helps ensure that they cannot interfere with one another in the event of a software failure in any partition. Also, applications are prevented from consuming excessive amounts of resources, which could starve other applications of the resources that they require.

Partitioning is also a means of achieving better hardware utilization when software does not scale well across large numbers of processors. In such circumstances and where possible to do so, running multiple instances of an application in separate smaller partitions can provide better throughput than running a single large instance of the application. Thus, LPAR, in some cases, can be the most cost effective means of improving response time and throughput for business operations.

4. Support of applications requiring different operating system versions

Every business has some old application whose origins are lost in antiquity, but whose function is critical to the business. Moving this application to a different operating environment, even a new version of the operating system, is fraught with risk. In addition, Linux operating system-based applications are increasingly finding their way into enterprise IT infrastructures.

LPAR provides an ideal way of accommodating the needs of various applications for different versions of operating systems. Starting with AIX 5L Version 5.1, different versions of AIX, as well as Linux, can exist within different LPARs in the same pSeries system. This enables a single system to have different versions of the operating system, and even completely different operating systems, installed to accommodate multiple application requirements. An LPAR can be created to test applications under new versions of the operating system prior to upgrading the production environment. Instead of having a separate server for this function, a minimum set of resources can be temporarily used to create a new LPAR where the tests are performed. When the partition is no longer needed, its resources can be incorporated back into other LPARs. One or more partitions running Linux can exist to facilitate the easy incorporation of this new environment, as its adoption within the organization moves forward.

5. Support of applications requiring different time zone settings

Many applications depend on the system time, which is set by the system administrator. Applications that support different regional operations usually run on separate instances of the operating system. Even if the applications are able to manage the different time zones themselves, it is still difficult to schedule system downtime for planned maintenance and upgrades without impacting regional operations. LPAR makes this issue easier to deal with.

Multiple regional workloads can be consolidated onto a single server. The different workloads can run in different LPARs, with different operating systems and different time and date settings. For example, workloads for operations based in London, Tokyo and New York can run in different LPARs on a single server. The evening batch workload, maintenance, or upgrade for the New York operation will not affect those of the London or Tokyo operations.

6. Minimization of planned system downtime

It should be apparent from the previous discussions that the flexibility inherent in LPAR greatly aids the scheduling and implementation of normal upgrade and system maintenance activities. All the preparatory activities involved in upgrading an application or even an operating system could be completed in a separate partition. The amount of time where business operations would be impacted could be reduced to a minimum.

7. Better utilization of scarce or expensive resources

Making commonly used peripherals such as tape drives, tape libraries, optical storage or high-performance communication adapter cards, available to individual partitions is simple and easy. Individual I/O slots can be assigned to an individual partition and then released for use by another partition. The ability to reassign I/O adapters and devices between partitions means that these can be purchased based on workload and usage needs rather than a “one per server” approach.

III. Quick Comparison of Vendor LPAR Approaches

The major UNIX vendors all claim to implement LPAR to some extent. However, the differences in implementation are critical to the realization of LPAR benefits. Technical discussions and explanations of LPAR implementations can be found in other white papers^{1, 2, 3}. The following table gives a quick overview of the different LPAR approaches from IBM, Sun Microsystems, and Hewlett-Packard.

¹ “IBM @server pSeries LPAR Technology” white paper by Edward Boden, Cesar Diniz Maciel, and Richard Cutler, IBM International Technical Support Organization, October 4th, 2001

² “Sun Solaris 9 Containers--Logical Partitions?” white paper by David Black, IBM pSeries Competitive Sales, May, 2002

³ “An Introduction to LPAR--Logical partitioning offers new systems management capabilities for pSeries administrators” white paper by Bill Buros, October, 2002

Comparison of LPAR Implementations From IBM, Sun, and HP⁴

	IBM	Sun	HP
Minimum # of CPUs per partition	1 CPU	Whole system board 2 or 4 CPUs	2 CPUs (4 recommended)
Smallest CPU increment to partition	1 CPU	One CPU if space available on system boards, else another system board	One CPU (4 recommended)
Minimum memory in a partition	256MB	All memory in a system board (2GB minimum ⁵)	All memory in a cell (2GB minimum, 4GB recommended)
Minimum number of I/O slots in a partition	Two PCI slots (can be reduced by using "built-in" network or boot disk connections)	All I/O slots on a system board (4 Sbus or 2 PCI)	All I/O slots on an internal I/O chassis (12 PCI slots)
I/O slot increment in a partition	One PCI slot	One Sbus or PCI if there is space on a system board	12 PCI slots if an I/O chassis can be attached to a cell
Partition incremental resources restricted by physical configuration	Yes (depending on model, some I/O slots must be in the same partition)	Yes, one or more system boards with ⁴ <ul style="list-style-type: none"> • 2 or 4 CPUs • 2 or 4GB memory • 4 or 6 I/O slots 	Yes, one or more "cells"; Each "cell" has system board with <ul style="list-style-type: none"> • 1-4 CPUs • 2-16GB memory
Separate boot adapter and disk required per partition	Yes	Yes	Yes
OS support per partition	AIX 5L v5.1 AIX 5L v5.2 Linux	Solaris 8 Solaris 9	HP-UX 11i
Dynamic reconfiguration of resources between partitions	Yes	Yes	No
Allocate/Deallocate I/O slots between partitions	Yes	No	No
Performance dependent on physical location of resources	No (however performance boost from memory affinity may be affected)	Yes	Yes
Partitioning available on mid-range or entry servers	Yes	No	No

⁴ Source of Hewlett-Packard and Sun Microsystems information "IBM @server pSeries LPAR Technology" white paper by Edward Boden, Cesar Diniz Maciel, and Richard Cutler, IBM International Technical Support Organization, October 4th, 2001

⁵ Dependent on the particular Sun computer model.

Overview of the advantages and limitations of each approach:

IBM

IBM's LPAR implementation allows a partition to be created with a minimum of:

- 1 processor
- 256MB of memory
- 1 boot disk
- 1 I/O slot for boot disk access
- 1 I/O slot for Ethernet access to the Hardware Management Console (HMC)

The I/O slot requirements can be satisfied by using "built-in" resources such as integrated Ethernet controllers and external SCSI connections. These resources can be physically located anywhere in the server (or external to it) so the ability to use any resource in any partition is uncompromised. LPAR is supported not only on IBM's high-end POWER4-based UNIX pSeries servers, but also on many of the POWER4-based mid-range and entry servers as well. IBM's LPAR implementation includes support of multiple versions of AIX as well as Linux. Dynamic LPAR requires AIX 5L v5.2.

Sun Microsystems

Sun provides their versions of LPAR based on physical subsets of resources. In Sun's case, the minimum physical building block consists of a system board, which includes:

- 2 or 4 processors
- Up to 4GB of memory
- 2 I/O buses which supports one or two I/O slots

Each Sun partition (or domain) is created with one or more system boards. However, the resources on one system board can only be used in one domain. It is not possible to configure CPU, memory, and I/O resources into different domains if they are located on the same system board. This does restrict one's flexibility in configuring and dynamically adjusting partitions. And Sun does not extend its domain capability to its "Value" class of servers, the V480 and V880.

The Sun domains can run different versions of the Solaris operating system, and the applications running in different domains are isolated and cannot interfere with one another. There is no support of Linux in Sun domains. Additionally, the Sun Enterprise systems do have single points of failure and a number of major shared components. Some of these components, depending on the nature of the failure, can disrupt the entire system. Failures on the main bus (called the Centerplane) and failures on the control board can cause all the domains to fail.

Hewlett-Packard

HP also provides their versions of LPAR based on physical subsets of resources. HP's LPAR approach is called nPartition. Every nPartition contains one or more cells and runs a separate instance of the HP-UX operating system. Each cell consists of a system board with

- 4 processors
- 2-16GB of memory
- connected to a single I/O chassis containing 12 I/O slots
- at least one I/O chassis must be connected to each nPartition and cannot be shared between multiple nPartitions

HP recommends that each cell within an nPartition have the same number of active processors and the same amount of memory configured in 4GB blocks. There are other restrictions related to physical locations of the resources in order to maximize performance. Should resources need to be physically reconfigured, an HP Service Engineer must carry out the reconfiguration, except for I/O cards.

The only operating system capable of running in an nPartition is HP-UX 11i, which only supports static partitions. Hence, if cells are to be added or removed from an nPartition, those affected must be shut down and rebooted for the change to take effect. There is no support of Linux in HP nPartitions, and HP does not support nPartitions in their mid-range or entry UNIX servers.

HP has said that in the future, they will provide a version of HP-UX that supports dynamic reconfiguration of nPartitions. However, there may still be requirements to physically move resources for optimum performance and availability, which would require system downtime.

IV. Customer Scenarios Where LPAR Is Especially Valuable

Fewer systems, less systems management, lower costs

A major retailer needed to replace their legacy retail management software with an effective supply chain management solution. To enable this, the company was looking for a new system that offered large processing power and the flexibility to support future developments. Because these applications run for nearly 24 hours every day, reliability is absolutely critical to the solution.

The system had to be designed to cope with over 1TB of data covering products from hundreds of stores. The ability to partition the servers and run multiple instances of UNIX with different workloads was very important. Scalability is crucial since the company will be adding more workloads and modules to their applications in the future.

The partitioned environment allowed the customer to run 8 AIX operating system images with just 2 servers - reducing floor space requirements and decreasing management costs. One key attribute that a partitioned environment gives is the flexibility to change the size of the images up or down depending on the actual workload requirements. The company is continuing to change its stores and number/type of products regularly, so what works

today is not necessarily what is required for tomorrow. The flexibility LPAR provides allows the system to adapt to changing requirements

The first server, which carries the main applications, has 4 partitions: 1 large database partition and 3 application partitions. The second server is used for a forecasting application together with high availability and contains 4 partitions: 1 database, 1 application, 1 backup partition (for High Availability to server #1) and 1 spare LPAR for test and development.

Running multiple application and database layers on a single server in a near 24x7 environment calls for extremely high levels of reliability to help ensure business operations continue uninterrupted. The pSeries solution is the ideal choice as it can deliver this capability and give the company confidence to operate their business without worrying about their IT infrastructure.

Smaller Companies Benefit Too!

A small Internet game company, located in Seoul, Korea, uses advanced and patented technology to offer a system and method for operating Internet games on communications networks. With over 60 employees, the company reports yearly revenues of approximately \$30 million.

They had two Sun E3500 servers running Solaris, two Intel® servers running Oracle 9i and three Intel “white box” servers running Linux and its own Web application. The number of people who use their game site is growing rapidly, and the company has been struggling with performance issues. The company sought a larger platform that would bring about more rapid Web application transactions. The company also wanted a more stable IT infrastructure, and it wanted to simplify server management.

IBM Global Services - Integrated Technology Services (ITS) suggested that the company consolidate all of its existing servers on to an IBM @server pSeries p690 server with two logical partitions (LPARs). One partition would act as the database server running AIX 5L, while the other would act as the Web application server running Linux. ITS ported the company’s existing database servers to AIX 5L and chose to maintain Linux as the platform for the Web solution. The robust pSeries hardware, with its outstanding performance and scalability has proven to be the ideal hardware platform for their needs.

The company expects to double the speed of its applications by running them on the pSeries server. Also, by using LPARs to consolidate their seven servers onto a single system, the company has considerably reduced its total cost of ownership. And, by maintaining Linux as the platform for its Web applications, the company can be sure that the application will be stable.

A Way To Provide Flexibility To Your Users

A utility company provides a service to their customers by selling processing capability to them. A major problem is that initial customer resource requests often lead to inaccurate server sizing. This in turn causes problems with applications when the server provided isn’t big enough. Using fixed configuration servers risks either offering their clients too much capacity, or too little capacity when the customer’s demands increase. There is no

way to adjust except by an expensive hardware replacement with a larger server.

Partitioned pSeries servers offer a fine solution to this issue. LPAR allows the customer to achieve operational flexibility. Initially the clients are offered relatively small LPARs (1 CPU, 1GB of memory and I/O) at a fixed price. Prices are provided for additional processors, memory and I/O. Actual LPAR allocation changes frequently due to changing client requirements and charges can be assessed in an appropriate and easily agreed on fashion.

The net result is that the utility can price client processing requests quickly and any unplanned requirements can be handled on the LPAR platform easily and efficiently. The 1 CPU, 256MB of memory, 2 I/O slot granularity that AIX LPAR offers is especially useful.

Capacity Upgrade on Demand (CUoD), offered on larger pSeries systems reduces initial hardware investment and provides even more flexibility. The spare processors provided under CUoD may be quickly activated, and more hardware obtained if requirements exceed even the installed capacity.

The customer now has a computing resource facility that they can assign and bill out easily and quickly to clients. They have secured a number of contracts to run applications in separate partitions while being able to adapt the configuration easily to changing client requirements. They can also support client needs for the support of Linux and different versions of AIX without having to purchase separate servers.

V. Conclusion

LPAR is not only a useful IT concept, but it is paying off in the real world. IBM's implementation not only provides significant advantages over other vendors, but LPAR is available on mid-range and entry servers as well. IBM LPAR also supports Linux in addition to AIX for even greater flexibility. LPAR is definitely a capability that can help businesses low costs and improve efficiencies.



Copyright IBM Corporation 2002

IBM Corporation
Marketing Communications
Server Group
Route 100
Somers, New York 10589

Published in the United States of America
11-02
All Rights Reserved

This publication was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this publication in other countries. The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM's future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, the e-business logo, AIX, AIX 5L, @server, POWER4, pSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both.

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

Linux is a registered trademark of Linus Torvalds.

Intel is a registered trademark of Intel Corporation in the United States and/or other countries.

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

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply.

Photographs show engineering and design models. Changes may be incorporated in production models.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

The IBM home page on the Internet can be found at www.ibm.com

The pSeries home page on the Internet can be found at www.ibm.com/servers/eserver/pseries