sageza

Strategic Snapshot

The Value of PowerVM Workload Partitions

New Virtualization Options in IBM AIX v6.1

By Clay Ryder

The Sageza Group, Inc.

February 2008

The Sageza Group, Inc.
32108 Alvarado Blvd #354
Union City, CA 94587
510·675·0700

sageza.com
info@sageza.com

The Value of PowerVM Workload Partitions

ABSTRACT

Many organizations have embraced virtualization to improve IT utilization and reduce the expenses associated with equipment acquisition, installation, and operation. While traditional virtualization or partitioning schemes have improved IT resource utilization, reducing the number of physical servers has not reduced the number of server operating system (OS) images requiring administration and maintenance. If anything, virtualization has encouraged growth in the number of servers that support the application workloads in organizations. There is an opportunity for IT to reduce this administrative overhead to become more streamlined and cost-efficient while continuing to provide the levels of service on which organizations have become dependent.

IBM AIX 6.1, through its support for Workload Partitions (WPARs), enables organizations to rethink the way they deploy multiple workloads on a single server. While traditional approaches such as virtualization using logical partitioning provide OS isolation and independence, for many workloads, this degree of isolation exceeds the user's need and results in unnecessary administrative and operational overhead.

WPARs offer IT managers a more cost-effective yet secured approach that meets the needs of many organizations. WPARs differ from other partitioning or virtualization schemes in that they partition server resources by the workload and share access to a single OS image. This is in contrast with the more common approach of creating a discrete operating system image to support each virtual server. By reducing the number of OS images required, the level of server software maintenance and other related IT administrative and management activity can be decreased while maintaining streamlined operational management and reduced need for physical resources.

WPARs increase resource utilization from the typical 5-20% average, reduce partition creation and teardown times, and reduce the number of OS instances and associated system management workload. WPARs provide standard application environments, support mobility and templates as well as cloning, and have automated policy-based resource and workload management through the WPAR manager. Consolidating with WPARs saves floor space and reduces the power consumption and expense associated with servers and air conditioning in the data center while maintaining the one-app/one-server deployment paradigm.

In this paper, we examine the flexibility that WPARs offer IT professionals in their virtualized UNIX server environments. In particular, we review how WPARs are different from other partitioning technologies and how WPARs complement existing environments. We discuss the capabilities and practical uses of WPARs in sample scenarios and articulate the ways in which WPARs provide an alternative to other partitioning schemes. Through AIX 6.1 and its support for WPARs and PowerVM Live Application Mobility, IT managers have greater flexibility in server configuration and can select the best approach to meet the user organization's needs while also simplifying the operational and cost efficiency of the IT environment.

The Value of PowerVM Workload Partitions

TABLE OF CONTENTS

What Are Workload Partitions?

One factor that has helped virtualization find acceptance in many organizations is that beyond the initial setup and deployment of the virtual server environment, the remaining IT tasks are largely unchanged. While this familiarity implies a small learning curve, it may not necessarily provide the most cost-effective IT environment. Since each virtual server has a discrete operating system (OS), tasks such as image and configuration, patch management and upgrades, and system backup must be performed for each virtual server. As a result, the degree of system configuration and maintenance remains relatively unchanged compared to a series of standalone servers, which over time can represent a substantial administrative cost.

A Workload Partition (WPAR) is a software partitioning scheme that is provided by AIX 6.1 and differs from hardware partitioning solutions such as Logical Partitions (LPARs). WPARs are not dependent on hardware features. AIX 6.1 and WPARs are supported on all servers based on the POWER4, POWER5, POWER5+, and POWER6 processors. WPARs are contained within LPARs and can be moved between various LPARs. As we will see, WPARs address different needs than other partitioning schemes and offer flexibility in virtualization.

WPARs are bundled with AIX version 6.1, the latest version of IBM's UNIX operating system. WPARs allow a more cost-effective yet simplified, consolidated, and secured approach to supporting multiple workloads on a single server. This new approach delivers efficient application deployment, while streamlining operational management and peacefully coexisting with current partitioning and virtualization solutions.

The value of the WPAR approach is that it reduces the number of OS images whereas LPARs focus on maintaining complete isolation of system resources, which requires discrete OS images for each partition. LPARs provide a discrete OS environment allowing a single server to run different types or versions of an OS. For example, Linux, i5/OS, AIX5.3, and AIX 6.1 can run concurrently on a single Power Systems server. WPARs can be used in workload consolidation initiatives; however, their primary aim is to reduce administrative and management expense by reducing the number of OS images required to support application workloads.

Capabilities of WPARs

WPARs provide a software-based partitioning mechanism for creating virtualized operating-system environments to manage multiple workloads. Since support for WPARs is provided by IBM AIX 6.1 and is not dependent on the underlying hardware, WPARs can be created on current technology POWER6-based servers as well as older servers based on the POWER4 and POWER5 processors. The underlying principle of these virtual environments is the sharing of an OS image within a POWER processor LPAR. There are two WPAR types: System WPARs and Application WPARs.

System WPAR

A System WPAR is similar to a typical AIX environment (with certain restrictions) in that each has a dedicated writable file system, though it can share some of the global environment's file system in read-only mode. A system WPAR has its own runtime resources, and when started it behaves like a standard AIX environment by creating an *init* process, which in turn spawns other processes and daemons as required. Consequently, the daemons that support networking capability, remote login to the WPAR, scheduling, and so forth are initiated to create the operating system environment that has its own System Resource Control.

Each System WPAR has its own unique set of users, groups, and network interface addresses. The users and groups defined within a system WPAR are completely independent from those defined at the global environment level. In particular, the *root* user of the WPAR is limited in that its superuser privileges exist solely within the WPAR; it has no such privileges or access to the global environment. For example, a DB administrator within a System WPAR hosting a DB server can be given root privilege, but its scope is limited to the WPAR, not the global environment.

Application WPAR

An Application WPAR is a lightweight resource designed to support a specific application. It contains one or more applications that can be started through a single AIX command line. The command line is passed as an argument to the *wparexec* command that will create and invoke the Application WPAR. This acts as a shell that spawns the application and supports the processes related to the application. The Application WPAR inherits the security context of the Global Instance (users and groups) and the Global Instance administrator is responsible for managing the Application WPARs; there is no separate *root* user for an Application WPAR. The partition will remain in effect until the command line that was passed in terminates, which results in the WPAR itself being terminated.

Processes executing within an application WPAR can only see other processes within the same WPAR. Inter-process communication by application software is limited by the boundaries of the WPAR. An Application WPAR shares the file system of the global environment, as it does not have a dedicated storage/file system. It can be configured to receive the application's file system resources from disks that are owned by the host AIX image including those from an NFS server. The partition can run daemons; however, it cannot execute system service daemons. It is therefore not possible to perform a remote login to an Application WPAR or remotely execute an action into the WPAR. Although Application WPARs do not have system services such as telnet that would support a remote connection, they do have unique IP addresses by which external connections to application services can be made.

Other Considerations

Both types of WPARs offer many benefits that permit maximum flexibility in partition configuration. Any WPAR provides discrete administration and security isolation from other partitions on the physical server. WPARs also offer administrative and operational savings as their key tenet is the reduction in the number of OS images deployed to support application workloads. Besides reducing the initial number of images to deploy, all subsequent management tasks such as patch maintenance, configuration, etcetera are similarly reduced.

Resource management is based on the AIX Workload Manager, a fair-share scheduler that allows organizations to control resource allocation based upon business priorities. Managed resources include CPU, memory, threads, and processes.

Live Application Mobility

If the business operations require more resources applied to a given workload yet the LPAR or physical server has insufficient spare resources to meet need, WPARs can be moved amongst LPARs on the same or different servers without affecting application availability. There are two methods by which WPARs can be relocated: PowerVM Live Partition Mobility and PowerVM Live Application Mobility.

Power VM Live Partition Mobility, which is supported by POWER6 and PowerVM Enterprise, allows for the relocation of an entire LPAR from one physical system to another

including any WPARs running inside of the LPAR. Live Partition Mobility requires that both LPAR have access to the same storage device (the application data will remain in place) and all I/O must be accessed through Virtual I/O Servers (VIOS). This relocation occurs without significant interruption to end users.

Power VM Live Application Mobility complements existing virtualization and/or relocation schemes as it allows active relocation (hot-migration) of WPARs and the applications they host without stopping application execution. In the case of Live Application Mobility, only the Workload Partition and the applications running inside of the WPAR are relocated between AIX instances running inside of LPARs; the AIX operating system is not relocated. This is done invisibly to the end user with the only notable difference being a somewhat longer response time while the application is migrating. Live Application Mobility is supported on any server supported by AIX 6 including servers based on POWER4, POWER5, and POWER6 processors. Application Mobility helps improve application availability, facilitates consolidation, allows for easy OS upgrades, and can be automated by the WPAR Manager using policy-based rules. Application Mobility is provided by AIX 6.1 and the WPAR Manager.

How WPARs Complement and Differ from LPARs

At first glance, it might seem that an organization would need to choose either workload or logical partitioning for a given server; however, this is not the case. WPARs are non-disruptive to LPARs and are actually complementary. The differences between the two approaches are not matters of technical incompatibility, but rather that each addresses a specific and different set of needs.

LPARs: Rock-Solid Isolation

LPARs are designed to provide rock-solid isolation of images and application execution. As such, this partitioning scheme is ideal for workloads that must achieve the highest level of isolation from all other virtual environments running on the same server. A guiding principal is that system events occurring outside of the LPAR will not affect operation within the LPAR. If an application or OS image in another partition should crash, its effect will be limited to that partition alone. Likewise, applications in a test environment that might inadvertently lock up system resources will not affect the normal operation of applications in other partitions.

To achieve this high degree of isolation, each LPAR is treated as if it were a standalone server. For example, each partition has a unique OS image and configuration. A separate software stack and dedicated system resources are also dedicated to an LPAR. As a result, each LPAR has a highly isolated environment protecting it from operational issues taking place elsewhere in the system. The principal difference between LPARs and WPARs is that LPARs can have any supported OS each with its own patches and tuning, whereas WPARs all have exactly the same base OS, i. e., the same version, patches, and tuning options. LPARs provide the most adaptable application run-time environment. However, this also means that each LPAR must be administered discretely, so that all system patching, logging, event monitoring, or application configuration required of a standalone server must be replicated for each LPAR deployed.

WPARs: Cost-Effective Isolation and Mobility

By deploying WPARs inside of LPARs, enterprises can reduce the ongoing costs of server software maintenance as well as reduce the number of administrative tasks associated with application deployment and maintenance. This is made possible through WPARs sharing an

OS image as opposed to each partition having its own, allowing reduced administration and maintenance overhead due to fewer OS images being required to support a given set of workloads. As many systems administration functions still rely on some degree of human intervention or activity, any reduction in the number of OS images to maintain offers a reduced exposure to misconfiguration or other human error.

WPARs provide organizations with the flexibility to deploy additional workloads while reducing the resource requirements commonly associated with partitioning. A shared OS image requires fewer system resources than multiple discrete images. In particular, the shared image approach can result in a smaller memory map dedicated to the OS as well as fewer system processes executing on a given physical server that are dedicated to overhead. Consider a common AIX configuration with 1GB of RAM dedicated to the OS. For fifteen LPARs, this would require 15GB of RAM, whereas fifteen WPARs could be configured using as little as 1GB or 2GB of RAM. Further, through Live Application Mobility, if a WPAR requires additional system resources that are not available on a given LPAR or server, the WPAR can be migrated to another LPAR or different server altogether without suspending the executing application(s) or interrupting user access.

Application workloads within a WPAR have equal access to all OS controlled resources, which can yield reduced isolation for each specific application when compared with a single application within an LPAR approach. As a result, WPARs do offer a lesser degree of isolation than LPARs; however, for many application workloads the isolation provided by WPARs is "good enough" to meet the customer need. In many cases, the reduced administrative expense is viewed as paramount over the enhanced operational isolation offered in other approaches.

As with any technological approach, there are advantages and disadvantages that must be considered in light of customer requirements and corporate best practices. The availability of WPARs offers organizations greater choice in their virtualization strategy and allows these decisions to be made on business or cost-effectiveness basis as dictated by risk/reward ratios, not just the limits of historic approaches to virtualization and partitioning.

Benefits and Usage Scenarios with WPARs

Organizations can derive several benefits from deploying WPARs as part of the AIX environment. Since the partition is a software construct, it is flexible and can be adjusted to meet the needs of the IT administrator over time. Once defined, a WPAR can be easily started from a command line either manually or in response to other system activities.

The AIX 6.1 Workload Partition Manager (WPM) provides the centralized management and resource optimization required to support application and system WPARs. Physical servers can be consolidated and deconsolidated dynamically in conjunction with Advanced POWER Virtualization capability afforded by the IBM Power Systems and AIX. Applications requiring less than one-tenth of a processor to support their workload can elect the WPAR approach to be part of a larger application consolidation into a Global LPAR that can deliver highly granular CPU and system resource utilization. This improved utilization offers enhanced ROI and the potential to delay future investment in additional resource.

To illustrate some of the benefits that WPARs afford, we asked two IT professionals about their experiences with WPARs to help illustrate resource savings they have achieved as well as new solutions they were able to develop using WPARs. In the remainder of this paper, we articulate some of the experiences of Henning Gammelmark, Systems Programmer with Bankdata Consulting, a Danish IT services organization focused on the banking industry, and Jez Wain, Systems Architect for Group Bull, a French IT systems vendor.

Reduced Administrative Expense

Organizations actively seek to not only reduce or delay their investment in new IT equipment but also to reduce the total cost of ownership (TCO) over the equipment's lifetime. Skilled management and administration of these servers is paramount to operational and financial success; however, for many organizations the need for IT administrative staff is poorly supported by corporate management.

Jez Wain relates a scenario from a study that Group Bull undertook for a large European telecommunications vendor. The IT environment needed to support 150 AIX servers or logical partitions. A traditional LPAR-only approach to server provisioning would have required using 1GB of RAM per AIX image for each of 150 LPARs and 1TB of disk space. Rule of thumb for administrative coverage is thirty LPARs per administrator; therefore, five FTE IT administrators would be required. The time to create and deploy the 150 LPARs was estimated at two man-months. Overall, the stated need was for five LPARs for each application, to provide online, remote, development, test, and other environments. All of these resources and expense would be just to support the OS; none of the resources for applications had yet to be applied.

By deploying WPARs as an alternative scenario, the potential savings in both human and equipments resources is striking. Through a 30-WPAR to 1-LPAR ratio, the same computing infrastructure could be delivered with five LPARs, 15GB of RAM, 75GB of disk storage, one administrator, and an install time of four man-days or less. This assumes that 3GB is reserved for each AIX image, which in reality is substantially higher than the amount required for smooth operation. Performance measurements between WPARs vs. a single AIX server in an LPAR had little discernible difference. Further, the teardown or reconfiguration time for a WPAR is matter of seconds and unless the file system is saved, the entire environment is wiped clean. This contrasts with the amount of effort it takes to configure and deploy an LPAR with a separate AIX instance.

Bankdata's Henning Gammelmark notes that using WPARs has saved time and effort while providing more flexibility and options in their WebSphere environment. Developers need test systems for many different WebSphere versions. Historically, this was supported by having discrete LPARs for each version, or by performing a backup/restore when a different version was required. This is time-consuming not only in doing backups/restores but also in maintaining the various versions as patches and updates are applied, which then requires restore and backup of each environment. By using WPARs, the OS is only patched once for the LPAR and is shared by each version of WebSphere deployed. In fact, Gammelmark notes that Bankdata is now able to support more test images than before, and they can be brought up or down in a matter of minutes, not hours.

As we have seen, since the WPARs share an OS image, TCO can be improved by reducing system administration costs due to the fewer number of OS images being maintained, and correspondingly fewer bug fixes, patches, and other updates to apply, as well as other maintenance-related activities. WPARs also provide performance optimization whereby applications or virtual servers can be scaled up or down depending upon the actual throughput and performance requirements. WPARs thus offer considerable reduction in resources required to support a given workload and sharply reduced provisioning time, a benefit welcome in any organization.

Administrative Delegation

As was mentioned earlier, a system WPAR maintains its own runtime resources and behaves much like a standard AIX environment with a unique set of users and groups. Since the *root*

user's superuser privileges exist solely within the WPAR, certain individuals could be given superuser privilege solely within the realm of a WPAR and not cause an undue security or operational integrity risk to the server as a whole. Administrative functions such as backup, user account maintenance for a specific application/function, application installation and configuration, and so forth could thus be undertaken by a subordinate IT professional or, in some cases, even general administrative personnel.

Normally such personnel would never be granted root access to a server, but for the purposes of a WPAR, it becomes feasible to delegate some administrative function safely away from senior IT professionals who hold the passwords to system-wide privileges. Jez Wain notes that this approach has resonated with many of Bull's customers as the WPAR owner has *root* access to application, but the administration of the server (LPAR) is now segregated from the administration of the application (WPAR). As a result, application administrative functions can be redeployed outside of IT to perhaps less costly human resources without placing the system as a whole at risk. The limited scope of the *root* user works to contain any accidental misconfigurations/changes in the environment to the defined limits of the WPAR. In some cases, it may become feasible to add more self-service attributes to deploying applications within WPARs. Preconfigured environments could be launched by general administrative personnel in a very straightforward fashion.

Flexibility in Deployment and Redeployment

The nature of computing workloads is dynamic and fluid with a rhythm that changes throughout the day. Some workloads also follow a peak that corresponds with longer periodic cycles such as the calendar month or quarter. In addition, at times it is necessary to shut down a server for planned outages and other maintenance operations. Given that WPARs can be relocated from LPAR to LPAR without interruption to application execution or its users, deploying applications within a WPAR can be very handy in these scenarios.

There are other scenarios as well where this mobility is beneficial. Bankdata's Gammelmark cites the ease with which he moved an updating DB/2 environment (that was contained within a WPAR) from one LPAR to another and then committed the update once the move was complete. Creating a DB/2 WPAR was a matter of minutes, whereas creating an LPAR and provisioning it with a new AIX instance would have taken much longer.

Similarly, Bull System's Jez Wain states that mobility can play a helpful role in testing and certification of application builds. Once an executing application passes the test environment, it can be migrated to the online production partition. This bars any last-minute changes to the underlying code or libraries from taking place as the application goes online in exactly the same state as in the testing environment.

Application mobility also benefits application responsiveness by relocating a workload to a different LPAR or physical server should the current host lack sufficient resources to maintain application service levels. Likewise, as workload demands diminish, the ability to consolidate active workloads onto fewer LPARs implies that some physical resources could be released, or even powered down altogether, if the IT workload/usage profile warrants.

Energy Savings and Consolidation

Energy consumption by IT has come under scrutiny as energy prices have risen sharply during the past few years. At the same time, the power consumption of ever more powerful servers has placed IT departments in the unenviable position of justifying further expenditures not only for new equipment but also for the energy to operate and cool it as well as the floor space on which to house it. Organizations may have historically dedicated one POWER-based server or a large LPAR to support each application, yet the workload often

requires only a fraction of the available processing power. This one-application/one-server approach is especially common in older installations with servers based on processors before POWER4 that did not support virtualization.

WPARs combine the advantages of both a scale-out and scale-up approach. Scaling out gives each workload its own blade or standalone server, which is exactly the same approach as a WPAR. The difference is that a blade or small standalone server cannot be driven at high levels (>50%) of CPU utilization, due to the capacity requirements to support infrequent peak loads, which in most cases allow servers to achieve average utilizations that are only in the high single digits. With WPARs, organizations can scale up multiple applications hosted on a single LPAR/physical server while still providing a segregated and dedicated application environment. By hosting multiple workloads together, it is possible to achieve a mix of workload peaks and troughs: the greater the variety of applications, the flatter the aggregate workload curve and the higher the safe utilization rate of the server.

Jez Wain of Bull Systems illustrates these efficiency improvements. Consider an environment built upon generic four-way x86 servers. Overall, these servers will draw ~800 watts, which will yield ~800 watts of heat. Cooling systems as a rule of thumb consume power at parity with the heat being dissipated; so 800 Watts are used by the air conditioning to cool each server, with a total energy cost per server of 1.6 kW. If each system achieves 20% CPU utilization, then four such servers could be consolidated onto a single server. The energy draw would be reduced from 6.4kW to 1.6kW for the same workload, a 75% energy savings. If server utilization were 10% as opposed to the very optimistic 20%, then energy savings would be ~87%: a very substantial reduction.

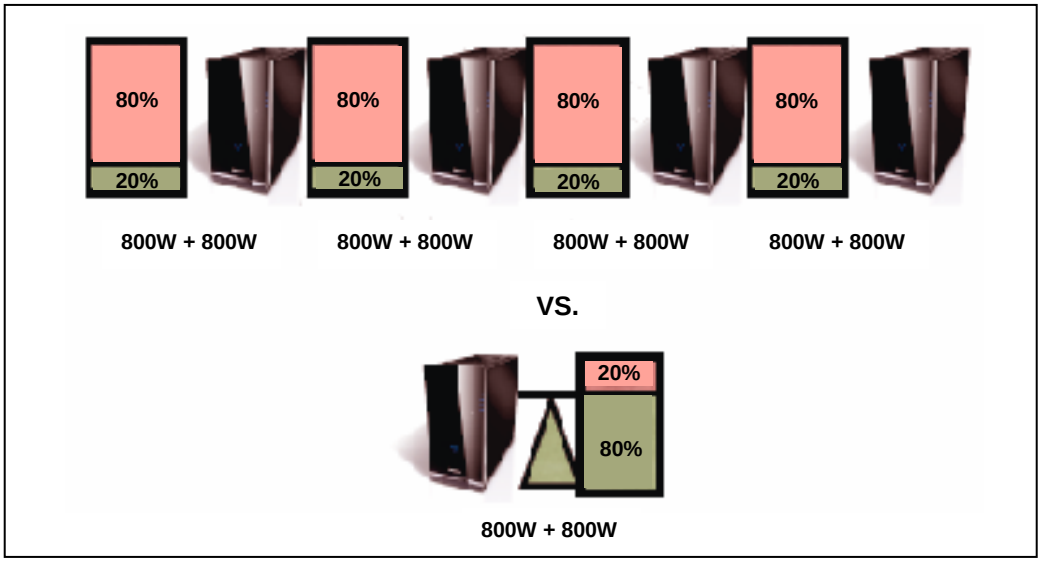

Figure 1: Potential Energy Savings through Consolidation

Additional energy savings can be achieved on POWER6 servers as Live Application Mobility can consolidate workloads during low utilization onto a reduced number of servers. Conversely, as demand increases, workloads can spread out or migrate onto a larger pool of servers. For a datacenter that experiences heavy workloads twelve hours a day, five days per week, this equates to only 35% of the time that servers must be in a peak load configuration. The remaining time offers opportunity for energy savings and overall cost savings.

Enhanced Application Availability

Although WPARs should not be considered a High Availability solution in the classic business continuity or unplanned outage sense, WPARs enable organizations to reduce

substantially the occurrence of planned application outages. Through the WPAR Manager, partitions can be relocated in response to a planned system outage, to improve application performance, or to increase overall LPAR/system utilization. In the case of planned outage, the benefits are obvious, as applications can remain available during otherwise "dark" hours. Application mobility can be set to operate on an automated basis, which automates much of the workload management based upon specified configuration rules.

Henning Gammelmark of BankData envisages the use of WPARs and Application Mobility to improve application uptime in their production environment. He notes that BankData has some systems running in a HACMP-like cluster, i. e., one node is active, and the other is a passive node that is completely idle. To affect a rollover to the passive node, the present configuration incurs application downtime. By changing the configuration to make use of WPARs with Live Application Mobility, Gammelmark envisions doing away with an active/passive scenario in favor of WPARs, which can be simply migrated as needed to accommodate planned system downtime. He estimates that moving an active database, which would take about an hour in the current configuration, could be achieved in less than ten minutes through WPAR mobility.

Resource Optimization and Standardizing the Application Environment

Just as an AIX image can be shared across multiple WPARs, so too can application code and libraries. This is especially handy when supporting multiple instances of standardized application environments. Jez Wain of Bull Systems cites an example:

If several instances of a Web server are required, one copy of the APACHE Web server software can be installed in an LPAR. WPARs created within the LPAR will execute the same APACHE libraries, as AIX will recognize that each WPAR is using the same version of the binaries and therefore will load them to memory only once. Each Web server in each WPAR has a discrete operating environment as well as dedicated disk space, but shares the Web server read-only code. This improves memory optimization through the elimination of duplicate code while reducing maintenance associated with the Web server code as only the single copy needs maintenance regardless of the number of actual Web servers being executed.

Additionally, this approach can have the effect of raising the granularity of management from a process level to the WPAR level. With discrete applications executing within an LPAR, it is necessary to review each application at the process level to assess whether sufficient resources are available or if adequate performance is being achieved. At the WPAR level, if the application is underperforming, the priority of the WPAR can be adjusted, or additional resources added to the WPAR, and the effects filter down to each process within the WPAR. Since the goal of most process tuning is to achieve better application performance, making the changes at the WPAR level simplifies the management process.

What It All Means

Throughout this report, we have highlighted many of the ways in which WPARs can reduce the complexity of administrating and maintaining operating system images while improving the overall operational efficiency of servers. The use of WPARs enables organizations to deploy additional application workloads with lower administrative overhead while delivering higher overall utilization levels and lower resource requirements.

WPARs are a complementary virtualization and resource sharing solution that co-exists with logical partitioning schemes to offer IT managers greater flexibility in their server configuration and administration. This permits IT professionals to select the most appropriate partitioning technology for each customer's need as well as being able to raise the overall utilization of their server resources.

The "one application on one server" model, while initially simple to deploy, is a costly and inefficient way in to scale IT resources within an organization. As needs scale, the complexity and waste of resources such as energy, hardware, cooling, administration, etcetera quickly increases until the computing solution devolves into expensive and unmanageable server sprawl. Virtualization technology addresses many of these issues while providing stable isolated execution environments; however, in many circumstances, the proliferation of virtual servers entails as much administrative support as standalone servers. In such environments, the IT design goal can often be extended beyond simply reducing the number of physical servers to reducing the total number of operating system images to be maintained. This is where WPARs can play an important role in the data center.

Live Application Mobility provides an easy-to-deploy mechanism by which active applications can be moved either manually or automatically from one WPAR to another without suspending their execution. While virtually eliminating planned application downtime, this mobility can lower the cost of server administration as it can sharply reduce the requirement that general IT maintenance take place during "off hours."

For organizations that have made virtualization a part of their IT strategy, managing a rapidly growing pool of applications is a very real concern. These organizations struggle to meet industry best practices while cost-effectively delivering applications and complying with various regulatory initiatives. WPARs offer greater virtualization and operational efficiency while allowing the deployment of additional workloads at lower administrative overhead due to the shared approach to the virtualized environment. As part of a larger virtualization strategy across the data center, WPARs offer organizations greater server utilization, the need for fewer physical servers, enhanced administrative efficiencies, reduced complexity, and decreased demands for energy and cooling, while at the same time delivering enhanced flexibility and quicker time-to-market for new IT services.

WPARs are an example of the continued investment of IBM in its UNIX platform hardware and software. This commitment to R&D is manifest in the way AIX 6.1 is able to deliver significant flexibility and user choice in server virtualization solutions at optimum cost effectiveness. IBM has a unique breadth of technology, products, and service that continues to evolve to meet the needs of organizations that have invested in a UNIX server strategy and solution set. WPARs build upon the value delivered by LPARs to offer organizations additional flexibility and choice. The combination of WPARs and LPARs provide IBM clients with unique virtualization options and business value across the entire Power Systems product line. IBM's continued investment in AIX and POWER processor technology benefits customers through enhanced flexibility in workload virtualization while simultaneously addressing the operational cost and complexity of IT environments charged with supporting the growing number of applications workloads found in most any organization.