IBM AIX 5L Version 5.2 for POWER, the On Demand UNIX Operating System

July 13, 2004

Abstract

This paper is intended to introduce IBM @server® pSeries®, @server p5, RS/6000® and @server i5 clients, IBM Business Partners, Sales, marketing and technical teams to the AIX 5L™ Version 5.2 for POWER™ operating system (OS). All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

The AIX 5L Value Proposition

AIX 5L V5.2 is a 64-bit UNIX® operating system for enterprise computing from IBM. AIX 5L is designed to exploit the advanced 64-bit IBM POWER processor-based systems and software architectures which include uniprocessor low-end servers; high-density, rack-mounted Web servers; clustered highly-available and scalable servers; symmetric multiprocessors (SMPs) for high-performance; and highly scalable parallel processors. AIX 5L V5.2 (with the 5200-04 Recommended Maintenance Package (APAR IY56722), or later) supports the new POWER5™ @server p5 and @server i5 systems further enhancing AIX 5L scalability and performance.

With the AIX 5L V5.2 operating system, systems are well suited for the mission-critical requirements of enterprise computing environments. The AIX 5L OS is based upon on demand computing and provides support for large-scale SMPs, cluster and parallel processing, dynamic logical partitioning and autonomic computing. AIX 5L offers clients a wide range of availability, scalability, connectivity and security features to fully leverage their investment in POWER processor-based systems.

AIX 5L Version 5.2 delivers innovative technologies for on demand computing including:

Dynamic reconfiguration for processors, memory and I/O
Dynamic Capacity Upgrade on Demand (CUoD) via dynamic logical partitioning
Workload management
Enterprise level distributed systems management
Enterprise storage management
Cluster Systems Management (CSM) for monitoring and administering multiple AIX 5L and Linux® systems from a single point-of-control (separately orderable)
Industry standard security and directory features
Enterprise level network features and connectivity
Development and performance tools and utilities

AIX 5L also supports IBM plans to establish Linux as a common application development environment for all IBM @server platforms. A key AIX 5L design objective is to continue to build and maintain a strong Linux affinity within the AIX 5L operating system.

Core Features of the AIX 5L Operating System

AIX 5L offers an impressive range of core features built around industry standards for systems management, security, network connectivity, enterprise storage, scalability and high performance computing (HPC).

Support for UNIX and Industry Standards

AIX 5L is an enterprise operating system environment based on UNIX (branded X/Open UNIX 98 Server and Workstation) and formal industry group standards (ANSI/POSIX, IEEE, ISO and FIFS). The operating system provides a wide variety of standard compliant features and software services that enhance application portability and heterogeneous system interoperability with systems such as Linux, Windows®, OS/2®, Solaris, z/VM®, z/OS®, i5/OS™, AS/400® and HP-UX.

A strong commitment to open standards such as UNIX, Linux, LDAP, SNMP, Java™ and XML is integral to on demand computing. AIX 5L brings a wide range of open standards that enable clients to manage new, open infrastructure solutions inside logical partitions and storage management frameworks on the pSeries systems. Open standards enable common software development of applications and middleware that can be easily ported across heterogeneous systems. POSIX Realtime signal support enhances application portability. pSeries, @server p5, and @server i5 systems provide significant investment in delivering Linux operating system support both native to the hardware and the Linux affinity feature in AIX 5L.

Linux Affinity

AIX 5L Linux affinity provides an integrated Linux application environment within the base AIX 5L operating system with the objective to support "compile and go" enablement for Linux applications developed on Intel® Architecture systems. It enables faster and less costly deployment of multi-platform, integrated solutions across AIX 5L and Linux platforms. Through CSM, AIX 5L supports both Linux and AIX 5L heterogeneous distributed cluster systems management. AIX 5L also offers a Web-based systems management client specifically for Linux.

IBM provides the AIX® Toolbox for Linux Applications, a collection of Open Source and GNU software commonly found with Linux distributions. Over 360 separate development tools, server applications and Linux utilities are available in both Open Source code and binary RPM (Red Hat Package Manager)format for AIX 5L. The AIX Toolbox for Linux Applications is available at no charge on CD-ROM or from the following IBM Web site:
http://www.**ibm.com**/servers/aix/products/aixos/linux/

Java

The Java language has rapidly gained importance as a standard object-oriented platform-neutral programming language since its advent in late 1995. For IBM, Java is a worldwide, cross-divisional development effort with significant ongoing research contributions from multiple IBM research and development centers. This collaboration produces one IBM enterprise class, quality Java code base for use by all IBM platforms. A single code base ensures consistent functional

implementation and perpetuates common fixes, function, testing, service, and improvements across all platforms.

Systems Management

For managing and maintaining the operating system configuration, AIX 5L offers several choices, from the command line and menu-driven System Management Interface Tool (SMIT) to a variety of remote management solutions such as Wireless System Management. The Web-based System Manager, a highly scalable, multiple host view of the administration environment provides secure host management with optional Secure Select Layer (SSL)-security. Other key management services include Reliable Scalable Cluster Technology (RSCT) infrastructure, Network Install Manger (NIM) and the Common Information Model (CIM) standard infrastructure for managed resources.

For managing and maintaining large numbers of systems in clustered configurations, AIX 5L offers two choices: Cluster Systems Management for AIX 5L (CSM) or the Parallel System Support Programs for AIX (PSSP).

SMIT Management

The System Management Interface Tool (SMIT) is a simple, yet powerful, tool that assists system administrators in performing system management tasks from menu-driven interfaces. All major system management tasks are presented in the main SMIT menu, providing a single entry point from which to start a task. Using fast paths takes the administrator directly to task menus or dialogs and eliminates the need to remember complex command syntax, valid parameter values and custom shell path names. As new or improved system administrator features and functions are added in AIX 5L, SMIT will be enhanced to provide a single consistent interface covering all system management functions.

Web System Manager

The AIX 5L Web-based System Manager is a comprehensive suite of system management tools that exploits familiar end user interaction concepts. It utilizes a management console capable of administering multiple AIX 5L hosts from AIX, Windows or Linux remote clients. With Web-based System Manager, any operating system and platform for which a Java 1.3 enabled browser is available can be used to manage a pSeries system. In addition, Web-based System Manager supports dynamic monitoring of system events through its integration of AIX Resource Monitoring and Control (RMC). The RMC subsystem is a powerful and flexible monitoring system that provides dynamic status updates, e-mail notification, and unattended responses to system events.

Security

AIX 5L is an enterprise level open UNIX operating system with a wide variety of services and features to support a versatile workload environment. To minimize the number of possible security exposures, the system administrator can use AIX 5L facilities to identify necessary services and shut down unused services. AIX 5L OS hardening is a global philosophy of system security that focuses not only on detection, but also on prevention. Its key features involve

removing unnecessary services from the base operating system, restricting user access to the system, enforcing password restrictions and controlling user and group rights.

Certified at the Common Criteria Security Certification (CAPP/EAL4) level 4, AIX 5L V5.2 makes integral use of strong, industry standard security and directory technologies to provide the highly secure environment required by government and business clients. The CAPP/EAL4 level of certification is also available on the POWER4™ firmware/hypervisor, extending the security certification to AIX 5L V5.2 logical partitions. AIX 5L expands these technologies with integrated support for Pluggable Authentication Module (XSSO/PAM), user-based PKI certificates, Enterprise Identity Mapping (EIM), Kerberos V5, BIND V9, SNMP V3, Mobile IPv6, Wireless Access Protocol (WAP) v1.1, OpenSSH v3.4 and multiple cryptographic libraries. These features are in addition to continued support for the IBM LDAP DirectoryServer, ICSA Certified IPSec/VPN secure networking and various Java security technology offerings.

Network Connectivity

AIX 5L supports a wide range of network availability, manageability, security, and performance features. AIX 5L provides an integrated suite of performance optimized TPC/IP protocols. The integration starts at the network kernel device drivers and extends through the protocol layer stack (network interface, network, transport, and application). AIX 5L incorporates hardware assist, a large degree of kernel caching with minimal data movement through the Internet (IP), and user datagram (UDP) and transmission control (TCP) protocols to deliver leading edge network performance. IP over Fibre Channel and EtherChannel (IEEE 802.3ad) support enhances network configuration flexibility. Dynamic tracking of Fibre Channel devices further provides enhanced reliability.

TCP/IP Quality of Service (QoS) is a family of evolving Internet standards under development by the Internet Engineering Task Force (IETF) to provide the means for preferential treatment to certain types of IP traffic. AIX 5L provides a wide range of QoS services including the ability to classify outbound traffic into distinct classes of service and to announce and establish resource reservations as requested by client applications. The AIX 5L QoS facility enables administrators to deploy and enforce network policies governing the use of network bandwidth. Key QoS features include differentiated services, traffic policing, in-profile and out-of-profile packet marking, traffic shaping, metering, integrated services for client and server applications, RSVP signaling, guaranteed service, controlled-load service and policy-based networking.

Internet Protocol Version 6 (IPv6) is the next generation of Internet protocol designed to expand the IP address space from 32 bits to128 bits, thus enabling a virtually unlimited number of global Internet IP addresses. The IPv6 enhanced addressing capability, along with new functions for end-to-end security, improved mobility support, simplified address configuration, and management make IPv6 a critical component in the evolution of e-business, especially for wireless Internet access through PDAs, pagers, mobile PCs, and wireless Internet business applications. AIX 5L has enhanced the migration path from IPv4 to IPv6 by moving IPv6 enhancements for IPSec, Quality of Service, automatic configuration, routing tables and mobility support into IPv4. AIX 5L also supports mobile devices using Mobile IPv6, thus providing those devices the ability to maintain the same network global address all over the world while changing locations.

AIX 5L provides an integrated, highly reliable set of Internet network server facilities including, domain name server (DNS), HyperText Transfer Protocol (HTTP) server, dynamic host configurations protocol (DHCP), virtual Private Network (VPN), open secure shell, SNMP and virtual private network support. These Internet server facilities can be configured into high availability server configurations using High Availability Cluster Multiprocessing for AIX 5L (HACMP™) to provide the fault resistance and high end scalability required by mission-critical Internet network servers to meet the needs of today's global Internet business environments.

Beginning with AIX 5L V5.2, clients can use the AIX 5L native Multiple Path I/O (MPIO) feature to define alternate paths to a device for failover purposes. *Failover* is a path-management algorithm that improves the reliability and availability of a device because the system automatically detects when one I/O path fails and re-routes I/O through an alternate path. Load balancing policies can be assigned to enhance storage bandwidth performance. All SCSI SCSD disk drives are automatically configured as MPIO devices. Other devices, such as, Fibre Channel storage devices can be supported, providing the device driver is compatible with the MPIO implementation in AIX 5L. Both IBM and various OEM FC storage devices are supported.

Enterprise Storage

A highly integrated, fault resistant, enterprise scalable multi-level storage facility is another core feature of the AIX 5L operating system. The AIX 5L enterprise storage facility provides standards-based support for current and emerging storage technologies including an integrated logical volume manager, journaling file system, and flexible support for Network Attached Storage (NAS) and Storage Area Network (SAN) environments.

The Logical Volume Manager (LVM) included in AIX 5L controls disk resources by mapping data between a simple, logical view of storage space and the actual physical disks. This logical view of the disk storage is provided to applications and users and is independent of the underlying physical disk structure. Data on logical volumes appears to be contiguous, but might not be contiguous on physical storage. This specific feature allows file systems, paging space, and other logical volumes to be resized or relocated, span multiple physical volumes, and have their contents replicated for greater flexibility and availability.

The Journaled File System 2 (JFS2) is the native file system included in AIX 5L. JFS2 uses database journaling techniques to maintain its structural consistency, thus preventing file system damage when the system is halted abnormally to support the entire set of UNIX file system semantics. The enhanced JFS2 included in AIX 5L is a next-generation enterprise scalable file system that provides fast system restart and improved file system metadata consistency through journaling. JFS2 introduces several key scalability and performance features including theoretical four petabyte file and file system capacity (currently tested and supported at 16 terabytes), extent-based allocation, efficient directory organization, concurrent I/O and dynamic file system object allocation. While tailored primarily for the high throughput and reliability requirements of enterprise servers, JFS2 is also applicable to client configurations where performance and reliability are desired.

Network Attached Storage devices are high-performance storage appliances that provide shared data and file serving to clients and servers on a LAN. Most NAS devices provide the capability to scale to terabytes of storage, as well as failover protection for high availability and data-redundancy policies for data protection. NAS devices usually offer a plug-and-play operation that exploits the existing LAN IP network infrastructure, thus providing easy installation, operation, and administration features. AIX 5L supports NAS storage devices using both file and block access methods. NFS and CIFS (Common Internet File System) network file access provide support for file level I/O operations that are application transparent whether accessing direct attached servers, LAN attached servers, or NAS devices. AIX 5L also supports the iSCSI protocol over IP to those NAS devices that support network accessed block I/O operations, thus enabling NAS support for general purpose storage applications.

A Storage Area Network is a high-speed network of direct connections between storage devices and servers usually connected by Fibre Channel (FC). The SAN can be viewed as an extension to the storage bus concept that enables storage devices and servers to be interconnected. The SAN uses similar elements as in Local Area Networks and Wide Area Networks such as routers, hubs, switches, directors and gateways. A SAN can be shared between servers and/or dedicated to one server. It can be local or extended over geographical distances. AIX 5L SAN support is used to connect shared storage arrays and tape libraries to multiple or clustered pSeries systems in order to improve both availability and performance.

SAN-attached storage provides many benefits to networked servers. Multiple servers can access multiple storage devices connected to the same network. Storage is independent of servers and applications. By providing access through multiple data paths, reliability, availability and serviceability are all improved. Storage processing is off-loaded from servers and moved onto a separate network, thus freeing up server capacity to focus on application workload. A single image of storage media in a centralized consolidated AIX 5L SAN simplifies management, enhances disaster protection and provides flexibility in scaling. AIX 5L SAN support also provides redundancy management through remote copy of data for disaster protection.

Finally, AIX 5L supports the industry standard SAN Management Common Host Bus Adapter (HBA) API which provides a standard and consistent interface for accessing information in the Fiber Channel Storage Area Network. This industry standard API enables use of SAN management tools on AIX 5L from Tivoli®, Veritas, EMC and other ISVs.

AIX 5L System Scalability and Performance

AIX 5L provides a highly integrated and tuned operating system environment that exploits the pSeries system design and POWER architecture processors to deliver exceptional performance across commercial and technical/scientific environments. POWER processor-based systems provide a highly scalable single-system AIX 5L OS image that will support up to 32-way SMP configurations configured with up to 1024GB of memory. Via special order, AIX 5L @server p5 and pSeries systems also support multi-system clustering configurations with thousands of cluster nodes. The broad range of vertical (scale-up) and horizontal (scale-out) scalability options natively supported by the AIX 5L operating system provides clients with flexible and efficient configurations to grow as required in a nearly seamless manner.

Symmetric Multiprocessing

AIX 5L exploits the SMP architecture of the POWER microprocessor to deliver high levels of computing performance. The operating system provides both 32-bit and 64-bit fully pre-emptible kernel implementations that can be selected at boot time. Both kernel implementations can efficiently support from 2- to 32-way multiprocessor configurations. AIX 5L also supports fine grain kernel locking with a highly optimized multi-thread processing design that exploits the M:N multi-threading system architecture. AIX 5L SMP enhancements combined with the POWER4+™ pSeries processor designs have resulted in outstanding SMP performance. AIX 5L V5.2 also supports the more powerful POWER5 processor used in @server p5 and @server i5 systems which further enhancing AIX 5L scalability and performance. AIX 5L V5.2 does not support the POWER5 simultaneous multi-threading capability or the Advanced POWER Virtualization option which are supported by AIX 5L V5.3.

AIX 5L V5.2 has been tuned and optimized to support up to a 32-way SMP image. Enhanced AIX 5L subsystems, I/O paths, and intrinsic system parameters enable large system scalability. Examples of this scalability include support for 100,000 open files per process, over one million open files per system, over 30,000 threads per process, one gigabyte of networking buffer pool, and over 130,000 each of message queues, semaphores, and shared memory regions.

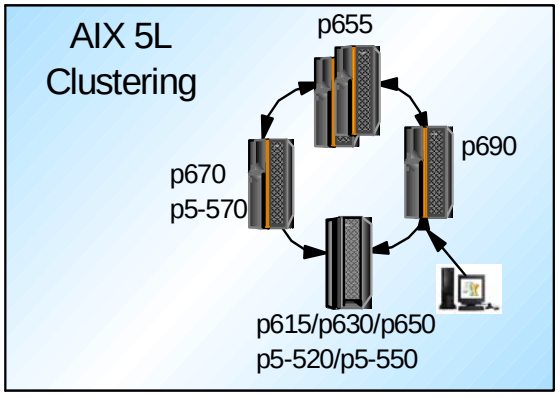

Clustering

AIX 5L supports a wide range of clustering software (CMS, HACMP and PSSP) for small LAN-based departmental clusters, to 32-system high-availability clusters, to massively parallel cluster systems containing thousands of servers or system nodes (special order).

A cluster is made up of two or more interconnected servers/nodes and/or logical partitions (LPARs) of servers, which are aggregated together and managed as a single, unified computing resource. The power of the cluster can be dedicated to a single, very large corporate application such as a DB2® Universal Database™, designed to scale incrementally and non-disruptively. Alternatively, the cluster can be shared among tens or hundreds of different enterprise applications including e-mail, Web services, enterprise resource management (ERP), client relationship

management (CRM), and supply chain management (SCM). Through clustering, enterprises can build cost effective IT infrastructures that can be easily managed from a single system administration point-of-control.

Clusters are designed to offer:
Large capacity data and transaction volumes, including support of mixed workloads
Horizontal and vertical scalability without downtime
Single point-of-control for distributed and clustered server management
Optimal use of IT resources with investment protection
24x7 access to data and applications
Business continuity in the event of disaster.

AIX 5L @server p5 and pSeries systems offer cluster solutions tailored to client requirements including products such as HACMP for AIX 5L or IBM @server Cluster 1600 with CSM or PSSP. IBM offers many software packages that will enhance the above cluster environments including:

- General Parallel File System (GPFS) for AIX 5L (high-performance cluster file system)
- Parallel Environment and the Parallel Engineering and Scientific Subroutine Library (parallel application development software, which uses the Message Passing Interface standard)
- LoadLeveler® (load balancing software designed to maximize throughput and increase job-turnaround in engineering and scientific computing centers)

Cluster System Management (CSM)

CSM for AIX 5L provides a powerful distributed systems management ability to cluster @server p5 and pSeries servers running the AIX 5L or Linux operating system and IBM @server xSeries®, IBM @server BladeCenter™ or IBM @server 325 servers running the Linux operating system and manage them from a single point-of-control. With this ability, the system administrator can run commands across multiple machines, gather responses, synchronize files, control hardware, manage power nodes on and off and monitor the health and status of systems in a cluster. CSM functions include automated set-up, hardware control, monitoring, and configuration file management. Administrators can manage individual servers or run automated responses to events across servers.

CSM is designed to exploit the advanced hardware management features of the @server p5, pSeries and xSeries servers. CSM offers a management environment that grows incrementally but does not increase management complexity as workload demands and the number of clustered servers increase. PSSP to CSM transition utilities are provided to ease the migration from existing PSSP SP environments to newer CSM configurations.

High Availability Cluster Multiprocessing for AIX 5L (HACMP)

HACMP for AIX 5L is designed to help ensure that applications and systems remain operational 24x7. HACMP can automatically detect system or network failures and provide the capability to recover system hardware, applications, data, and users. HACMP makes use of component

redundancy in a clustered computing environment to keep applications running, and to restart them on a backup system if necessary. HACMP clusters offer multiple data backup and recovery methods to meet disaster management requirements. HACMP also includes a High Availability Geographic Cluster (HACMP/XD) feature.

Up to 32 @server p5 and pSeries systems can be configured in a HACMP for AIX 5L cluster to meet complicated application availability and recovery needs. HACMP can be configured to react to hundreds of system events such as detection of software problems that are not severe enough to interrupt proper operation of the system. HACMP monitors, detects, and reacts to such failure events, allowing the system to stay available during random, unexpected software problems. As hardware has become more reliable, a significant cause of downtime is application failure. HACMP can be used in combination with LPARs to increase application availability. HACMP can monitor for application failure or reduced performance and restart the application on another LPAR. This flexibility enables the application to quickly return to a working state while the failed LPAR is diagnosed.

High Performance Computing (HPC)

AIX 5L offers a wide range of customizable features to optimize the pSeries system architecture and to maximize the performance of HPC applications and algorithms. Key HPC customizable features include: large page support, assignment of processor and memory affinity to specific processes, and support for optionally available optimized engineering/scientific libraries and compiler and tool support for parallel application development.

Large Page Support

AIX 5L and POWER4+/POWER5 processors support two virtual page sizes, a standard Power Architecture™ 4KB page size and an extended 16MB page size for both 32-bit and 64-bit applications. The use of 16MB large pages may improve the performance of both high-performance computing and business memory access intensive applications that use large amounts of virtual memory. An enhanced virtual buffer with the ability to map a larger virtual memory range improves application performance by reducing translation look-aside buffer misses. Large pages also improve memory pre-fetching by fetching 16MB pages, thus eliminating the need to restart pre-fetch operations on 4KB boundaries.

Processor and Memory Affinity

Processor and/or memory affinity configurations can improve the performance of memory-access intensive HPC jobs running on multiple multi-chip module (MCM) pSeries systems. Memory and processor affinity helps localize workloads by utilizing a subset of processors that share uniform access to a subset of system memory. The AIX 5L Workload Manager (WLM) allows system administrators to associate a set of processors with a WLM class so that jobs associated with that class only run on the dedicated processor set. The processor set can specify a single processor from the two processor chip which effectively dedicates all on chip cache and memory bandwidth to that processor. This configuration significantly improves sequential data intensive performance. AIX 5L also provides command and programming API support to bind processes to processors and further optimize performance.

AIX 5L Autonomic Computing

Autonomic computing is the term used to describe self-configuring, self-healing, self-optimizing, and self-protecting computer technologies. Autonomic computing represents intelligent, open systems that help manage themselves (self-managing). Servers are able to continuously tune themselves, adapt to unexpected conditions, help prevent and recover from failures and provide a safe environment for critical data and systems. AIX 5L is built on architectural foundations that encompass autonomic concepts. Specific examples of the autonomic computing features of AIX 5L include First Failure Data Capture and recovery, automatic system and I/O hang detection and recovery, self-optimizing disk management, dynamic partitioning for efficient resource utilization and the ability to automatically dial-up for service in anticipation of a system failure. Two key objectives of the autonomic computing architecture of AIX 5L are to automate systems management and to maximize system availability.

Automating Systems Management

The goal of automating systems management is to drive significant reduction in the costs associated with all aspects of managing pSeries systems. Elements of cost reduction include reducing the complexity of management tasks, decreasing the need for training and specialized skills, increasing the number of systems that a single administrator can manage, reducing administrative staff and eliminating or automating common management tasks.

Automatic system hang recovery with error detection and fix capabilities are key features of the automated system management of AIX 5L. AIX 5L can detect the condition that high priority processes are monopolizing system resources and prohibiting normal execution. AIX 5L offers system administrators a variety of customizable solutions to remedy the system hang condition.

I/O device hang condition detection and recovery due to I/O errors further illustrate the automated system management capability of AIX 5L. The system administrator can customize thresholds that, when reached, automatically initiate recovery procedures. POWER4/POWER5 systems can handle many uncorrectable data errors so that system operation continues until the data is actually used by a processor. With AIX 5L V5.2, if the data associated with an uncorrectable error is used only by a user process as opposed to the kernel or system firmware, the process using the data is terminated, thus avoiding system or partition reboot.

AIX 5L offers several additional tools for automated systems management. Alternate disk installation and migration provides system administrators the ability to migrate their systems to a new AIX 5L level dynamically while the system is still operating properly. Advantages of using alternate disk migration over a conventional migration are reduced downtime, quick recovery from migration failures and a high degree of flexibility and customization. Web systems management and SMIT systems management interfaces provide easy access to kernel tuning parameters including virtual memory manager, processor scheduler, Network File System (NFS), Input/Output, and network tunable parameters.

An automated dump analysis (ADA) facility for preliminary fault analysis of system crashes allows for early diagnosis of problems through auto retrieval of dump data which can be forwarded via e-mail for analysis. Dynamic application core dumps generate application core files without requiring application termination, thus helping to increase application availability and serviceability. Finally, Enterprise Identity Mapping (EIM), an IBM wide autonomic technology supported by all IBM @server systems, helps create a self-managing computing environment. EIM is a LDAP-based application that uses the LDAP database as a central repository of user mapping information. EIM can map a single network user sign-on identification and password across heterogeneous applications and systems throughout an enterprise.

Maximizing System Availability

The goal of maximizing system availability is to drive system performance to zero downtime. Maximizing system availability is accomplished with a variety of autonomic features. Processor deallocation and replacement is an example of the self-healing and self-optimizing capabilities of AIX 5L. The Dynamic Processor Deallocation function automatically and dynamically removes failing processors that reach a predetermined error threshold from a system image. The system automatically reassigns the workload to other available processors to avoid interruption. If the system must be rebooted, previously deallocated processors will not be included to avoid repetition of the error condition. Failing processors can be replaced during normal service to minimize system and application downtime.

If spare unlicensed processors are available on the systems via Capacity Upgrade on Demand (CUoD), they will automatically be enabled to replace the failing processor by the Dynamic Processor Sparing capability. When the failing processor is assigned to a dynamic logical partition, an inactive processor will be transparently activated and added to that dynamic LPAR, thus maintaining performance and improving system availability. The service processor records this action and notifies the system administrator of the condition. Computing operation continues with the processor deactivated allowing repair to be scheduled at a convenient time. When the failing processor is returned to service, the spare is returned to the inactive CUoD pool of resources.

An example of the self-protecting capability of AIX 5L is the Kerberos Authentication Server. Kerberos is a network authentication service that provides a means of verifying the identities of principals on physically insecure networks. Kerberos provides mutual authentication, data integrity, and privacy under the realistic assumption that network traffic is vulnerable to capture, examination, and substitution. Workload Manager underscores the self-optimizing nature of

AIX 5L by automatically balancing workload across systems. Workload Manager, in concert with the other autonomic computing features of AIX 5L, contributes to the attainment of maximum system availability.

On Demand Computing with AIX 5L

An on demand enterprise is one whose business processes are integrated end-to-end across the company and with key partners, suppliers and clients. An on demand business can respond with speed to any client demand, market opportunity or external threat. On demand computing refers to the IBM strategy to deliver computing resources, services, and solutions to clients on an 'as needed' basis across the business. A successful on demand computing environment will encompass five key fundamental technology strategies: open standards, virtualization, dynamic logical partitioning, CUoD and workload management.

A strong commitment to open standards such as UNIX, LDAP, SNMP, Java, XML and Linux marks one of most important elements of on demand computing. AIX 5L brings a wide range of open interfaces and services that for example, enable clients to manage new, open infrastructure solutions inside logical partitions and storage management frameworks on the pSeries systems.

System Virtualization

Virtualization promotes the ability to dynamically provision server, storage, and application resources according to changing demand and priorities. A virtualized on demand operating environment maximizes unused capacity by leveraging AIX 5L features such as dynamic logical partitioning, CUoD computing and workload management.

The key objectives of server virtualization are to help enable non-disruptive growth, cost effective management of variable business needs, efficient management of workload peaks and overall improvement in the efficiency of running a business. A key strength of the pSeries, @server p5 and i5 systems is their ability to concurrently execute multiple AIX 5L and Linux operating systems in up to 32 logical partitions. This capability can potentially consolidate a large number of AIX 5L and Linux servers on a single pSeries system.

Consolidation on a single system has many advantages. It considerably reduces floor space and cabling requirements. The number of outages and repair actions is reduced. And, perhaps most importantly, the cost for administration, management, and maintenance decreases. These advantages reduce the total cost of ownership (TCO) compared to a traditional server farm where applications run on a large number of separate servers. AIX 5L provides several key enhancements required for large scale server consolidation including the rapid creation of virtual servers for fast application deployment, support for dynamically increasing or decreasing LPAR resources based on run time demands and support for managing resource utilization so that resources are allocated as needed.

Dynamic Logical Partitioning (Dynamic LPAR)

Dynamic LPAR increases the flexibility of partitioned systems by providing system administrators with the ability to add and remove processors, real memory, and I/O slots from active LPAR partitions. This ability allows system administrators to assign hardware resources where they are most needed and to adjust to changing hardware requirements without impacting system availability. The addition and removal of hardware resources in a POWER4/POWER5 system are supported at the level of a single processor, a 256MB real memory block (1GB required for the first LPAR), and a single I/O slot. Under dynamic LPAR, system administrators can add and remove hardware resources through the Hardware System Console using a graphical user interface.

Static partitioning means that the resources (processors, memory and I/O) assigned to a partition cannot be reassigned without a partition reboot. Dynamic LPAR allows the administrator to reassign system resources "on the fly" without interrupting operations. This feature greatly improves operational flexibility and system availability. For example, an administrator may wish to have the bulk of their processors assigned to a partition that supports transaction processing during peak hours. During off-peak periods they may wish to reassign resources to manage back-office functions or complex queries. AIX 5L also supports moving infrequently used I/O devices such as DVD-ROMs or tape drives between partitions without causing a disruption. Resources can be removed from running partitions and used to create and boot a new logical partition.

Only partitions running AIX 5L V5.2 support dynamic LPAR. Other partitions running AIX 5L V5.1 or Linux can coexist on the same server with partitions running AIX 5L V5.2, but those partitions are managed as static LPARs and must be rebooted in order to add or remove resources. AIX 5L provides an API to control and monitor dynamic LPAR configurations. Applications or middleware programs can set events that will notify the program when dynamic LPAR resource changes occur. Programs can then take appropriate action based on resource changes such as enforcement for license restrictions, performance optimization, or user notification. Programs can also use the dynamic LPAR APIs to dynamically add or remove resources to the partition.

Capacity Upgrade on Demand

CUoD allows AIX 5L clients to purchase and enable extra hardware resource capacity initially unused by the system. When additional processors or memory are needed, the system administrator can dynamically enable the use of these extra resources through dynamic LPAR services without having to reboot the system or LPAR. With dynamic LPAR, CUoD provides fast non-disruptive upgrades that enhance flexibility and ease of use for variable workload management.

Capacity Upgrade on Demand provides clients with a clear economic benefit. Clients have the option to pay for additional pSeries processor and memory capacity only when needed, thus lowering the total cost of ownership. Key-based licensing allows administrators to quickly activate additional processor and memory capacity. CUoD is marketed in many configurations such as,

permanent, temporary and on a trial basis to meet varying business environments and requirements.

Workload Management

The AIX 5L Workload Manager (WLM) provides system administrators a policy-based method for managing the resources used by applications and users. WLM delivers automated resource administration for multiple applications running on a single server. This capability helps to ensure that critical applications are not impacted by the resource requirements of less critical jobs. WLM helps deliver the benefits of server consolidation and centralized systems administration.

Workload Manager provides the same types of functions within an AIX 5L LPAR as it does in a single SMP server. WLM is dynamic LPAR-aware and makes the appropriate adjustments when LPAR resources are added or removed, however, it currently does not control the addition or removal of LPAR resources.

Other key WLM features include:

- Support for use of either Web-based SM graphical interface or SMIT menus to create profiles (job classes) and manage job assignment
- Creation of profiles to allocate processor, memory, and disk I/O resource usage with minimum and maximum limits
- Assignment of jobs to profiles based on user/group ownership, application name or tag, process type, or manual assignment
- Automatic event notification for system administrators with ability to define responsive actions
- Flexible policy definition for processors associated with a WLM class

Conclusion

AIX 5L V5.2 is a premier UNIX operating system from IBM that delivers compelling value for clients in the on demand world. AIX 5L supports multiple industry standards to create an open and flexible environment including affinity for Linux applications. AIX 5L offers a rich feature set for systems management, security, network connectivity, enterprise storage, scalability and high performance computing. As the trend towards self-managing systems accelerates, AIX 5L is well positioned to help clients derive the maximum value from their pSeries servers.

The AIX 5L operating system is closely aligned with the IBM vision for on demand computing. Through system virtualization, dynamic logical partitioning, security, Capacity Upgrade on Demand and workload management, AIX 5L is a key enabler for clients embracing the on demand business model. IBM plans to continue an aggressive investment program in AIX 5L software development for the foreseeable future. Clients can be confident that AIX 5L will offer innovative new capabilities in the future to create value and respond to their on demand business needs.

Additional References

For additional information on AIX 5L and a discussion of IBM AIX 5L strategy and products, the following sources are recommended:

- The AIX 5L external on-line documentation Web site at:
  http://publib.boulder.**ibm.com**/cgi-bin/ds_form

- The AIX 5L Version 5.2 documentation index:
  http://publib16.boulder.**ibm.com**/pseries/en_US/infocenter/base/aix52.htm

- Information on AIX 5L development environment on IBM alphaWorks® Web site:
  http://www.alphaworks.**ibm.com**/

© IBM Corporation 2004

IBM Corporation
Integrated Marketing Communications
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
July 2004
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document 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 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, the AIX 5L logo, @server, AIX, AIX 5L, alphaWorks, AS/400, DB2, DB2 Universal Database, HACMP, i5/OS, LoadLeveler, OS/2, POWER, Power Architecture, POWER4, POWER4+, POWER5, pSeries, RS/6000, Tivoli, xSeries, z/OS and z/VM are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of U.S. trademarks owned by IBM may be found at http://www.**ibm.com**/legal/copytrade.shtml.

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Windows is a registered trademark of the Microsoft Corporation.

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

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. 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.

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 or other public sources. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided “AS IS” and no warranties or guarantees are expressed or implied by IBM.

The IBM home page on the Internet can be found at http://www.**ibm.com**.

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