

# AIX 5L: ESSENTIAL KNOWLEDGE GUIDE FOR POWER USERS



## TOPICS

|   |    |
|---|----|
| 1. Introduction to AIX 5L.....  | 3  |
| 2. Taking Advantage of Key New Features in AIX 5L V5.3.....   | 5  |
| 3. Growing Incrementally with Capacity Upgrade on Demand (CUoD) and Reserve Capacity On Demand.....         | 12 |
| 4. Achieving Higher Performance Levels with Simultaneous Multi-Threading .....                              | 15 |
| 5. Creating Logical Partitions with AIX 5L V5.2 .....   | 18 |
| 6. Creating Micro-Partitioning with AIX 5L V5.3 .....   | 24 |
| 7. Sharing Resources Between Partitions .....   | 27 |
| 8. Leveraging Virtualization of I/O and Network Resources with AIX 5L V5.3 .....                            | 31 |
| 9. Using the New Advanced Accounting Feature of AIX 5L V5.3 .....   | 33 |
| 10. Implementing Policy-Based Automatic Downloads of Updates with Service Update Management Assistant ..... | 37 |
| 11. Recompiling Linux Applications to Run in a Native AIX 5L Environment.....                               | 39 |
| 12. Additional Resources.....   | 43 |

## 1. INTRODUCTION TO AIX 5L

IBM AIX 5L™ for POWER™ Version 5.3 offers new levels of innovative and self-management technologies on advanced POWER5™ systems. AIX 5L is emerging as the preferred platform for UNIX® users and independent software vendors (ISVs). AIX 5L delivers industrial-strength UNIX reliability, availability and security, and offers flexible system administration and ease of integration with Linux®. With the optional IBM Advanced POWER Virtualization feature with advanced Micro-Partitioning™ and AIX 5L V5.3, eServer p5 systems deliver new levels of performance for the on demand world.

AIX 5L V5.3 continues to exploit the current 64-bit POWER5™ systems and software architecture. This AIX 5L release supports new POWER5 hardware systems, advanced virtualization, and simultaneous multi-threading POWER5 processors for improved performance and system utilization. AIX 5L V5.3 offers enhancements in scalability, performance, system management tools and improved system security. Additional features for enterprise distributed storage, application affinity and enablement, and networking (including NFSv4 support) illustrate how AIX 5L V5.3 delivers innovation and value for IBM clients.

The Advanced POWER Virtualization hardware feature supported by AIX 5L V5.3 on eServer p5 systems includes: Micro-Partitioning support for a single processor being shared by up to 10 dynamic logical partitions (LPARs); virtual Ethernet connections to provide high-speed connections between partitions; virtual SCSI disks that allow partitions to access storage without requiring a physical storage adapter and dedicated disk drive; and a new Partition Load Manager feature

---

*AIX 5L delivers industrial-strength UNIX reliability, availability, and security and offers flexible system administration and ease of integration with Linux.®*

---

providing automated CPU and memory resource load balancing across partitions. Additional AIX 5L V5.3 innovations include Advanced Accounting Feature to provide for accurate billing of IT resource usage and Service Management Update Assistant to provide policy-based automatic downloads to clients at an enterprise level of AIX 5L maintenance fixes.

## 2. TAKING ADVANTAGE OF KEY NEW FEATURES IN AIX 5L V5.3

AIX 5L V5.3 provides many new features and functions that enhance the existing enterprise level of operating system scalability, network environment, reliability, availability, standards support and rich native development environment. AIX 5L V5.3 extends its leading-edge standards support with its design to conform to the Single UNIX Specification V3 (SUS V3). This section will discuss the key new AIX 5L V5.3 features across these areas.

### ADVANCED JOURNAL FILE SYSTEM AND LOGICAL VOLUME MANAGER

AIX 5L V5.3 Advanced Journal File System (JFS2) and Logical Volume Manager (LVM) contain the following enhancements. The JFS2 file system shrink function supports optimizing storage utilization by removing unused disk space from the file system environment. Administrators can dynamically add and delete disk space as needed to manage both the JFS2 and LVM environments in place, without the need to copy and reboot. JFS2 supports the NFSv4 extended Access Control List (ACL) attributes and fine-grain permissions enhancing security. The maximum JFS2 file size has been expanded to greater than a 16TB size. JFS2 also provides support for limiting file system space to a specified size (quota) for users/groups supporting a managed resource environment. JFS2 enhances performance by supporting large-block I/O transfers that are only limited by the transfer size of the disk (up to 16MB) extending the previous 128KB block size limit. The LVM extends the maximum number of disks supported in a scalable volume group from 128 to 1024 disks.

## NETWORKING AND SECURITY ENHANCEMENTS

AIX 5L V5.3 provides an enhanced Network Install Manager (NIM) service handler that supports a trusted authenticated model based on key encryption. The NIM service handler communications now offer a secure infrastructure that is not dependent on the use of remote shell server for remote installs. Instead, using a client daemon with OpenSSL optionally configured within the NIM environment, NIM can provide a trusted encrypted authentication. The NIM service handler also simplifies setting up the NIM host/client environment with a host-name query feature that will automatically define client and network configurations. Additional NIM features are provided that supports the creation, synchronization and failover between multiple NIM masters supporting a highly-available NIM environment. The configuration of Etherchannel and Virtual IP Address (VIPA) network interfaces is added to NIM. Finally, NIM provides a hard disk scrubbing option at install time.

The following new AIX 5L V5.3 networking features are supported. The Network File System V4 (NFSv4) protocol support, IETF RFC 3530 compliant, is integrated into AIX 5L V5.3 with continuing compatible support for NFS V2/3. NFSv4 running over the TCP transport improves NFS performance, scalability and security. The Kerberos V5 security mechanism is used for authentication for both NFSv4 and V3 protocols. The NFSv4 ACL model supports fine-grain permissions, which enhance data access control.

The AIX EtherChannel supports Ethernet bandwidth aggregation from several Ethernet adapters. The new dynamic membership in EtherChannel feature allows users to add or remove an adapter from a running

---

*The Network File System V4 (NFSv4) protocol support, IETF RFC 3530 compliant, is integrated into AIX 5L V5.3 with continuing compatible support for NFS V2/3.*

---

EtherChannel without having to first detach its interface. This is useful when hot-plugging adapters that belong to an EtherChannel. AIX 5L v5.3 adds the TCP/IP support over Fibre Channel interface (IP over FC) (RFC 2625). The IP over FC feature supports IP packets sent over a physical Fibre Channel connection enabling the Fibre Channel adapter to be used as a LAN device. In a SAN environment, this will support servers and storage systems to communicate and perform networking types of management services without additional LAN equipment and infrastructure. The AIX 5L iSCSI protocol driver is included, providing storage device access over gigabit Ethernet TCP/IP networks.

The AIX 5L V5.3 Multi-Path I/O (MPIO) route selection policies have been extended to include random, weighted random, lowest utilized, weighted round-robin, and hash-based algorithms, increasing the system administrator's control over MPIO reliability and performance. The 10 Gbps Ethernet-SR PCI-X adapter is supported, providing greater bandwidth performance. Wake on LAN technology is supported to remotely wake up machines. The following new communication protocols are supported to improve clients' networking options and support new standards: Streams Control Transmission Protocol (SCTP), a reliable transport layer protocol over IP (RFC 2960); Service Location Protocol client API (RFC 2608/2614); Dynamic Host Configuration Protocol (DHCP) for IPV6 host and client (RFC 3315); and Advanced Sockets API for IPV6 (RFC 3542). AIX 5L V5.3 has also enhanced the network traffic debugging capabilities in the tcpdump and libpcap library.

AIX 5L V5.3 provides the following Lightweight Directory Access Protocol (LDAP) enhancements. The LDAP client support for server-based authentication supports

server-based authentication, rather than the previous client-based model enhancing security and interoperability. The LDAP client can connect to the LDAP server via Kerberos instead of having to rely on a hard-code password or stash file enhancing security. The LDAP client supports administrator specification of the LDAP server permitting administrators to set up different user communities on specific servers to provide load balancing across the LDAP environment.

### RELIABILITY, AVAILABILITY AND SERVICEABILITY

The Service Update Management Assistant (SUMA) is an AIX 5L V5.3 feature that provides administrators the ability to set up enterprise-wide AIX 5L maintenance databases that can support a variety of fix and maintenance levels. Clients can set up policies that automatically download AIX 5L maintenance data from the IBM eServer Support Fix Central Web site and then manage AIX 5L maintenance in an enterprise concentric model based on business requirements. Generic policies can automatically retrieve fixes in categories without administration action. Categories can include critical fixes, fixes based on a fileset or component or security fixes. Administrators are alerted when downloads occur and advised which fixes require reboot.

AIX 5L V5.3 provides the following reliability, availability and serviceability (RAS) enhancements. Improved handling of spurious interrupts makes the system less likely to hang as a result of receiving thousands of errors from a device (for example, temporary errors on a Fibre Channel link). The spurious errors could flood the error log and cause secondary errors. This change will allow the error log subsystem to avoid logging multiple duplicate errors, resulting in greater system reliability. Enhanced First

---

*Clients can set up policies that automatically download AIX 5L maintenance data from the IBM eServer Support Fix Central Web-site.*

---

Failure Data Capture (FFDC) causes key failure resolution data to be written to non-volatile RAM (NVRAM). This is useful in cases where the system dump is not successful. Corefile compression and other enhancements allow the user to specify the directory that core files will be written to and to compress the core files to about a 1 to 5 compression ratio. The SNAP subsystem is enhanced to provide a standard interface for third-party applications, to capture information that is specific to that application in the SNAP dump. Application developers can use this API to allow SNAP to gather relevant application debugging information. Error log hardening reduces the possibility of the error log being corrupted by an unexpected system failure. DVD media are supported as a dump device, increasing dump performance and reducing the number of dump media.

## ADDITIONAL ENHANCEMENTS AND TOOLS SUPPORT

AIX 5L V5.3 supports the POSIX real-time application interfaces (SUS V3 Realtime Option Group) including: barriers, spin locks, clocks, thread options, memory lock, priority scheduling, message passing, semaphores, timers, advisory information, asynchronous IO, fsync, mapped files, memory protection, real-time signals, shared memory, synchronized IO, thread priority inheritance/protection/scheduling, CPU time, and timeouts.

The AIX 5L V5.3 virtual memory pager is improved to release paging space blocks earlier, thus improving performance when removing a large number of pages. Support for a new memory allocation algorithm (Watson malloc) can reduce memory fragmentation and improve massively multi-threaded application performance. High-performance computing (HPC) environments may enhance performance with the support of a new memory

---

*Support for a new memory allocation algorithm (Watson malloc) can reduce memory fragmentation and improve massively multi-threaded application performance.*

---

algorithm that can efficiently page out large blocks of memory. New features added to the malloc debugging facility and the dbx symbolic debugger aid in the diagnosis of memory allocation problems. The dbx application debugging capabilities are further enhanced to include facilities for examining process information, file descriptors, kernel threads and core files, and new features for suspending breakpoints and managing loaded modules.

AIX 5L V5.3 provides new functionality in base commands, including tree-based process listing, time-stamped shell histories, cron logging controls, recursive file searching, restricted Korn shells, tar recursion control, new date formats, and vi back tagging, as well as other new capabilities for the make, find, fuser, restore, man, tar, at, nohup, and cron commands. The awk, ed, grep, head and vi commands can now process lines of at least 8192 characters.

The following AIX 5L V5.3 performance tools and libraries have been enhanced. The PMAPI (hardware performance counters application programming interface) library now supports threaded applications running in M to N mode. The perfstat library, and the vmstat, iostat, sar, topas, curt and splat tools have been enhanced to support Micro-Partitioning and simultaneous multi-threading. Filemon, netpmon and pprof have been enhanced to support Micro-Partitioning. The iostat, vmstat and sar commands now detect and tolerate dynamic configuration changes. Using the -d option of the sar command, average queue size, and service and wait time can now be monitored for selected disk types. The gprof tool now supports multi-process and multi-threaded applications. Trace now supports single process and thread tracing. Lparstat is a new tool that displays partition configuration information and supports monitoring of a set of partition

level performance metrics. Mpstat is a new tool that can be used to monitor a large set of detailed performance metrics at the logical processor level. Procmon is a new eclipse-based process monitoring tool, showing a fully configurable sorted list of processes, and offering access to selected performance tools to be applied to processes from the list.

Performance Toolbox (PTX) V3.1 provides extended support for the Advanced POWER Virtualization feature including the Partition Load Manager (PLM). On PLM-equipped systems, PTX can monitor special load/utilization averages, active processor counts and memory allocations that control the operation of the Load Manager. This provides a dynamic visualization of system performance, as PLM controls and adjusts resources to the system workload. For the Advanced POWER Virtualization feature, the PTX instrumentation can collect over three dozen new performance metrics related to simultaneous multi-threading and shared processor pool environments. These metrics include processor allocations (logical, physical, virtual counts, and capacities), memory allocations and capacities, and defined partition entitlements. Partitions can be viewed simultaneously in relationship to one another, for an overall view of system utilization. Changes in partition allocations are dynamically discovered and displayed for monitoring, recording, and playback. The Performance Aide V3.1 agent fully supports the collection and recording of the new PLM and virtualization metrics identified above.

---

*For the Advanced POWER Virtualization feature, the PTX instrumentation can collect over three dozen new performance metrics related to simultaneous multi-threading and shared processor pool environments.*

---

### 3. GROWING INCREMENTALLY WITH CAPACITY UPGRADE ON DEMAND (CUOD) AND RESERVE CAPACITY ON DEMAND

Capacity on Demand (CoD) adds operational and configuration flexibility for IBM eServer pSeries and p5 systems. Available for a fee, CoD allows additional resources to be added as they are needed on selected models. Processors and memory can be brought online to meet increasing workload demands. If the system is configured for dynamic LPAR, this can be accomplished without impacting operations. These features have significant value for users who want to upgrade without disruption, enhance their system RAS characteristics, or simply grow with a finer level of granularity. The CoD features allow a system to be manufactured (or upgraded in your locale) with inactive resources, such as processors and memory. The hardware is delivered with these features built in, ready to be activated when needed. The features can also be installed in the client's location. CoD can be turned on by ordering the type of activation desired, at which time the user is provided an activation code to enable the resource.

---

*Processors and memory can be brought online to meet increasing workload demands.*

---

An additional benefit of the CoD offering is Dynamic Processor Sparing. For environments with CoD, Dynamic Processor Sparing allows inactive processors to act as dynamic spares. An inactive processor is activated if a failing processor reaches a predetermined error threshold, thus helping to maintain performance and improve system availability. Starting with AIX 5L V5.2, this capability was offered on pSeries servers with CUoD processors to help minimize the impact to server performance caused by a failed processor. Dynamic Processor Sparing will happen dynamically and automatically when using dynamic

logical partitioning and the failing processor is detected prior to failure. If not detected prior to failure or not using dynamic LPAR, a reboot of the system or partition activates an alternate processor from the inactive spares. The user can then re-establish required performance levels without waiting for parts to arrive on site. Dynamic Processor Sparing does not require the purchase of an activation code; it requires only that the system have inactive CUoD processors.

## TYPES OF COD

After a system with the CoD features is delivered, it can be activated in the following ways:

- Capacity Upgrade on Demand for processors
- Capacity Upgrade on Demand for memory
- On/Off Capacity on Demand for processors and memory
- Reserve Capacity on Demand for processors
- Trial Capacity on Demand for processors or memory

**CUoD for Processors** and **CUoD for Memory** are activations which permanently enable the resource for future use.

**On/Off CoD** enables users to temporarily activate processors and memory. This offering provides a system manager an interface to manage the activation and deactivation of resources. A monitor which resides on the server logs the usage activity and periodically reports this information back to IBM. A usage bill is then generated based on the selected quantity and duration of processor and memory resources. Billing increments are in Processor Day and 1GB Memory Days.

**Reserve CoD** represents an “autonomic” way to activate temporary capacity. Reserve CoD enables the user

to place a quantity of inactive processors into the server's shared processor pool, which then become available to the pool's resource manager. When the server recognizes that the number of base (purchased/active) processors assigned across uncapped partitions have been 100% utilized, and at least 10% of an additional processor is needed, then a processor day (good for a 24-hour period) is charged against the Reserve CoD account balance. Another Processor Day will be charged for each additional processor put into use, based on the 10% utilization rule. After a 24-hour period elapses, and there is no longer a need for the additional performance, no processor days will be charged until the next performance spike.

## TRIAL COD

Trial Capacity on Demand (Trial CoD) is a function delivered with selected pSeries and eServer p5 servers supporting CoD. Those servers with inactive CUoD processors or memory will be capable of using a one-time, no-cost activation for a maximum period of 30 consecutive days. This enhancement allows for benchmarking of CoD resources, or can be used to provide immediate access to inactive resources when a purchase is pending.

---

*Reserve CoD enables the user to place a quantity of inactive processors into the server's shared processor pool which then become available to the pool's resource manager.*

---

## 4. ACHIEVING HIGHER PERFORMANCE LEVELS WITH SIMULTANEOUS MULTI-THREADING

Simultaneous multi-threading is a processor design that allows a POWER5 processor to issue instructions from multiple threads, each cycle using multiple on-chip thread contexts. That is, a single physical processor is now able to simultaneously dispatch instructions from more than one hardware thread context. The result is more saturated utilization of the processor resources, which can lead to a boost in overall throughput.

In the UNIX world, developers are familiar with process level parallelism (multi-tasking) and thread level parallelism (multi-threads). The introduction of simultaneous multi-threading is essentially a means for going from thread level parallelism into instruction level parallelism. Superscalar processors were designed with the capability to support multiple issue pipelines; however, in reality the instruction sequence of a normal single thread of execution rarely allows for full exploitation of the processor. Refer to “single thread” diagram in Figure #1 for a graphical explanation of this situation. The POWER5 simultaneous multi-threading design takes advantage of the superscalar nature to further saturate the processor resources with multiple hardware threads. There are two hardware threads per processor for the POWER5 design; see “Simultaneous multi-threading” in Figure #1.

---

*The POWER5 simultaneous multi-threading design takes advantage of the superscalar nature to further saturate the processor resources with multiple hardware threads.*

---

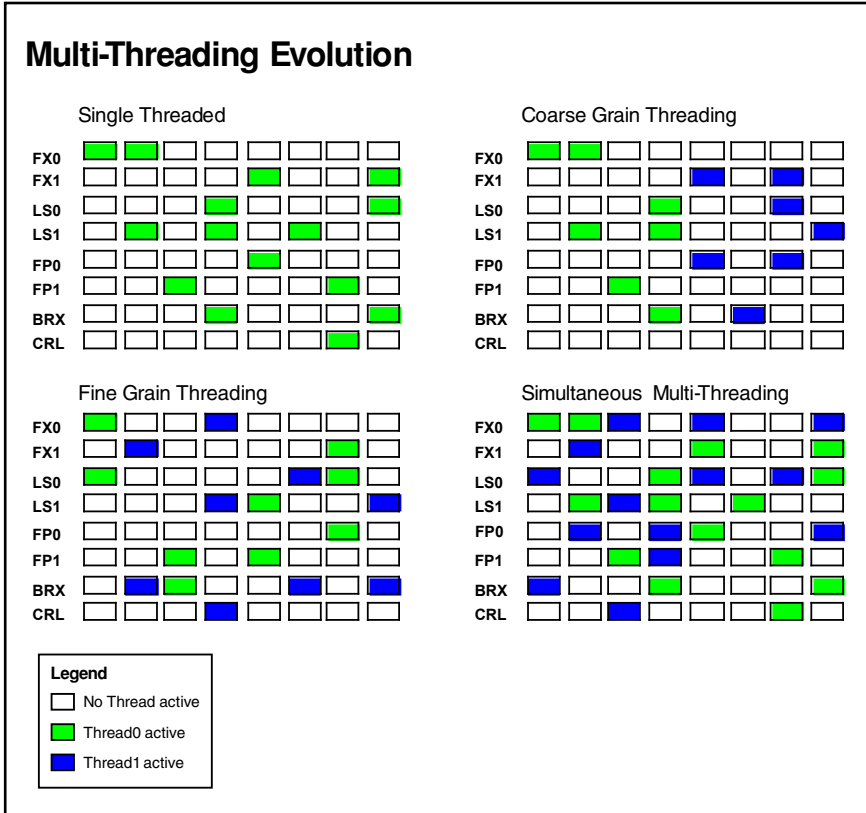


Figure 1. From Single Thread to Simultaneous Multi-threading

Based on the discussion above, it is clear that the performance gain of simultaneous multi-threading depends on the workload. The basic concept is that no single execution thread can fully saturate the processor, so it is preferred to have multiple applications generating multiple threads at the same time. The real benefit of simultaneous multi-threading comes in scheduling together threads with different processor resource utilization. For example, running a floating-point intensive thread with a fixed-point intensive thread allows for better saturation of both floating and fixed-point units.

In general, simultaneous multi-threading performance gain should be more significant for throughput oriented workloads.

It is expected to be used primarily in commercial environments, where the speed of an individual transaction

is not as important as the total number of transactions that can be performed. Simultaneous multi-threading is expected to increase the throughput of workloads with large or frequently changing working sets, such as database servers and Web servers.

Figure 2 below shows a POWER5 processor with two simultaneous multi-threading threads in details. Note that while the simultaneous multi-threading mode is on, AIX 5L V5.3 sees a dual core design as a 4-way processor chip.

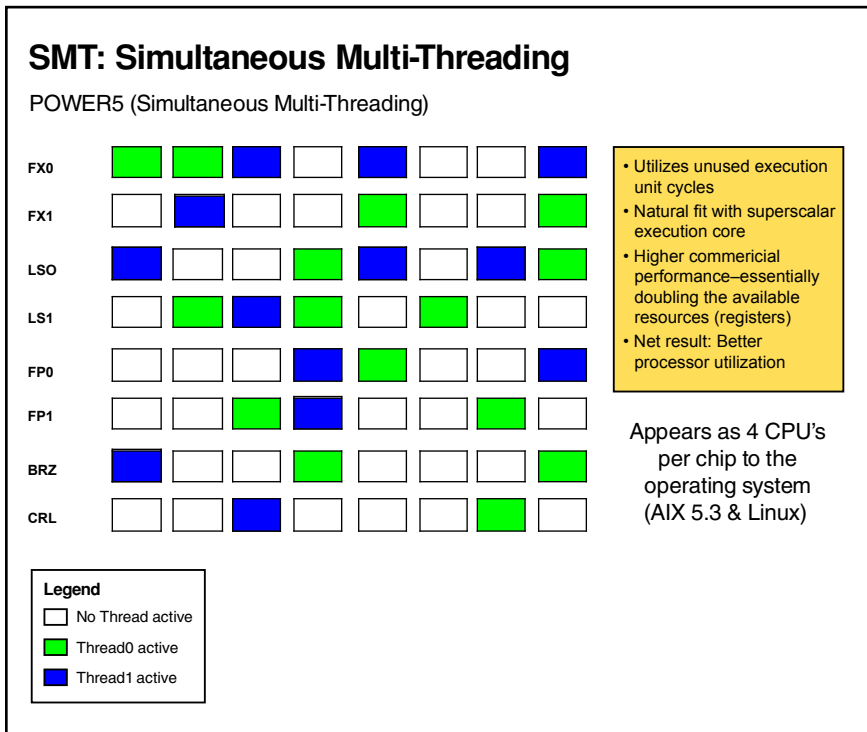


Figure 2. SMT details and summary

## 5. CREATING LOGICAL PARTITIONS WITH AIX 5L V5.2

Dynamic logical partitioning on AIX 5L V5.2 has greatly expanded server capabilities. A server can be divided into multiple partitions. These partitions become the hardware base for the deployment of applications and workloads. The dynamic partitioning feature allows the allocation and reallocation of system resources to happen while applications are running. The latest AIX 5L V5.2 maintenance release has further extended dynamic LPAR support to the POWER5 processor-based hardware systems.

### LPAR AND AIX 5.2 DYNAMIC LPAR

Logical partitioning (LPAR) is the ability to divide a single server into multiple logical machines, or partitions, that run operating systems and applications independently from each other. Logical partitioning allows flexible use of system components, such as CPU processors, memory, and I/O from the entire pool of available hardware resources, to construct the partitions, each with its own set of system resources, grouped and configured for the business need. Dynamic LPAR allows the resource allocation and reallocation to be done dynamically, i.e., it is not necessary to shut down the partition before adding or removing resources to or from the partition. AIX 5L V5.2 is dynamic LPAR-capable, and will stay up when the system adds or removes CPU, memory and adapters. These resources are configured dynamically for a partition, without a need for a shutdown and reboot of the system or the affected partitions.

---

*Dynamic LPAR allows the resource allocation and reallocation to be done dynamically, i.e. it is not necessary to shut down the partition before adding or removing resources to or from the partition.*

---

## AIX 5L V 5.2 DYNAMIC LPAR RESOURCES

An AIX 5L V5.2 OS-based dynamic LPAR can consist of the following resource elements: individual processor, 256MB memory region, and I/O adapter slot. Multiple resources can be placed under the exclusive control of a given logical partition. Dynamic LPAR extends these capabilities by allowing this fine-grained resource allocation to occur not only when activating a logical partition, but also while the partitions are running. Individual processors, memory regions and I/O adapter slots can be released into a “free pool,” acquired from that free pool, or moved directly from one partition to another.

Users of the affinity partition configuration option, which allocates CPU and memory resources in fixed patterns based on multi-chip module (MCM) boundaries, should note that only the I/O adapter resources can be dynamically reconfigured while in that mode.

## CONFIGURATION AND MANAGEMENT OF LOGICAL PARTITIONS

The configuration and management of logical partitions is done through a Hardware Management Console (HMC) which interacts with the Hypervisor layer to accomplish the partition management tasks. The HMC is a dedicated machine that runs on Linux and Java™-based software components to provide a graphical user interface and a set of applications for creating, configuring and managing the server partitions. The following software functionalities are provided:

- Create and store partition profiles, which define the processor, memory, and I/O resources to be allocated to a partition.
- Start, stop, and reset a partition or system.

- Boot a partition or system by selecting a profile.
- Display status for systems and partitions.
- HMC code update, debug, and error logging.
- HMC problem determination and service support such as call-home and error log notification.

The HMC can also be remotely accessed from a PC through the use of Web-based System Manager software. The PC client for the Web-based System Manager may be downloaded from the server machine for PC installation, or directly installed from the HMC.

---

*When creating a partition, a profile is also created to specify the minimum and desired set of resources for that partition.*

---

When creating a partition, a profile is also created to specify the minimum and desired set of resources for that partition. Each partition can have multiple profiles defined, with varying sets of resources, though only one profile can be active at a time.

When creating a profile for allocating CPU and memory to a partition, the system administrator defines a minimum number of processors and a minimum amount of memory required for that partition. Additional values define the desired maximum configuration for the partition. When activating the partition, the HMC checks if the minimum required resources are available. If they are, the partition can be started. If more resources than the minimum required are available, they are included in the partition, up to the desired value specified in the profile. The resources allocated to a partition can be changed by selecting another profile when reactivating the partition.

## APPLICATION CONSIDERATIONS

Applications, in general, do not require changes to run in AIX 5L V5.2 logical partitions, including those participating in dynamic LPAR operations. A partition is equivalent to a physical machine for an application. However, ap-

plications that are “bound” to specific system resources may need to make some changes to accommodate or to adapt to dynamic LPAR requests.

There are two different approaches that an application can take to ensure “problem-free” running in dynamic LPAR environment: be dynamic LPAR-safe, or become dynamic LPAR-aware. If an application is dynamic LPAR-safe, it is prevented from failing in the case of a dynamic LPAR event. It can also block dynamic LPAR events from affecting any of the system resources it already has. On the other hand, an application that is dynamic LPAR-aware will try to react and adapt to the resource changes of a dynamic LPAR operation.

### DYNAMIC LPAR-SAFE

Dynamic LPAR operations in AIX 5L V5.2 are designed to be non-destructive. Therefore, most applications are dynamic LPAR-safe by default. They will run in the dynamic LPAR environment without requiring any changes because most applications are not programmed to any specific system resource. Applications generally rely on the operating system to manage the system resources, such as CPU and memory. The following two situations are the exceptions where an application may not be dynamic LPAR-safe.

- An application code has been optimized for a uniprocessor and a CPU has been added.
- An application’s use of data is indexed by the CPU number and a CPU is added or removed.

The following two situations are the exceptions where an application is DLPAR-safe, but it may fail the dynamic LPAR requests for removing resources:

- If an application uses `bindprocessor()` and CPUs are being removed.
- If an application uses `plock()`, or `shmget(SHM_PIN)` and the memory is being removed. The dynamic LPAR request will fail if there is not enough pinnable memory available in the system to accommodate the migration of the pinned pages. This is a system level restriction that occurs when >80% of system memory is pinned. The remedy for this problem is to reduce the level of pinned memory in the system.

## DLPAR-AWARE

Applications that should be aware of the number of on-line CPUs and memory will be good candidates for making it dynamic LPAR-aware or to adapt to dynamic LPAR events. There are two ways that a dynamic LPAR event may be integrated into an application: a script based approach or a signal and API approach. The scripts are designed to externally reconfigure the application. The signal method assumes that the application is coded to catch the signal, `SIGRECONFIG`, and that the signal handler will reconfigure the application. Please see reference 1 for details.

## BENEFITS OF LPAR AND DYNAMIC LPAR

As a summary, the benefits of LPAR and dynamic LPAR are listed below for quick references. For a high-end server that has many hardware resources, LPAR can provide many benefits:

- Consolidate servers.
- Run simultaneous production and test environments.
- Consolidate multiple versions of the same OS.
- Consolidate applications requiring different time zone settings.
- Isolate an application.
- Have flexible workload scenarios.

With dynamic LPAR, you have increased resource flexibility that can adapt to changing workload demands even more dynamically than with LPAR. Here are some ways you can benefit from dynamic LPAR:

- Improved RAS (reliability, availability and serviceability). A defective CPU can be dynamically removed. Also, dynamic LPAR combined with Capacity Upgrade on Demand enables dynamic CPU sparing, where a defective CPU is transparently replaced with a spare, inactive CPU.
- Upgrade PCI system boards (with hot-plug). The upgrade can be accomplished without a reboot and without affecting other partitions.
- Shift resources to a partition where they are most needed.
- Dynamically allocate resources enabled through CUoD.
- Install a new operating system on a small partition for testing. In order to keep an application up-to-date with new operating system levels, partitions can be created on the fly for testing new operating systems. Dynamic LPAR makes it easier to free resources from active partitions, so that new partitions can be created.
- Reassign infrequently used I/O devices, such as the CD-ROM or tape drives.
- Move processors or memory between partitions during peak usages, and then move them back again after the peak has passed.
- Maintain a free pool of resources (CPU, memory, I/O) so that a newly-created partition's resources can be assigned easily from the free pool.

Reference:

Application in an dynamic allocation environment on AIX 5L V5.2 – <http://www-106.ibm.com/developerworks/eserver/articles/dlpar/>

Dynamic LPAR in IBM – pSeries <http://www-1.ibm.com/servers/eserver/pseries/hardware/whitepapers/dlpar.html>

LPAR for the IBM pSeries p690 system – <http://www-1.ibm.com/servers/eserver/pseries/hardware/whitepapers/lpar.html>

Logical partition pSeries LPAR related whitepapers – <http://www-1.ibm.com/servers/eserver/pseries/lpar/wp.html>

How to manage Dynamic LPARs and Develop Dynamic LPAR-aware Apps [http://www-106.ibm.com/developerworks/eserver/articles/pu\\_duttaDLPAR.pdf](http://www-106.ibm.com/developerworks/eserver/articles/pu_duttaDLPAR.pdf)

## 6. CREATING MICRO-PARTITIONING WITH AIX 5L V5.3

The Hardware Management Console which provides the user interface for logical partitions has been enhanced to accept the following partition attributes required to define the characteristics of shared partitions:

- Minimum, desired, and maximum processor capacity
- Minimum, desired, and maximum number of virtual processors
- Capped or uncapped
- Variable capacity weight

Processor capacity attributes are specified in terms of processor units. One processor unit represents one percent of one physical processor. In other words, one hundred processor units is equivalent to one physical processor.

Partitions are configured to run within the shared processor “pool.” A partition may be defined with a processor capacity as small as 10 processor units with additional increments of 1/100th processor, up to the size of the pool. This minimum assignable capacity represents 1/10th of a physical processor. A maximum of 10 partitions may be started for each physical processor in the platform.

When a partition is started, the system chooses the partition’s entitled processor capacity from the specified capacity range. The value that is chosen represents a commitment of capacity that is reserved for the partition. This capacity cannot be used to start another shared partition, otherwise capacity could be overcommitted. Preference is given to the desired value, but these values cannot always be used, because there may not be

---

*A partition may be defined with a processor capacity as small as 10 processor units with additional increments of 1/100th processor, up to the size of the pool.*

---

enough unassigned capacity in the system. In that event, a different value is chosen, which must be greater than or equal to the minimum capacity attribute. Otherwise, the partition cannot be started.

The same basic process applies for selecting the number of online virtual processors, with the extra restriction that each virtual processor must be granted at least 10 processor units of entitlement. In this way, the entitled processor capacity may affect the number of virtual processors that are automatically brought online by the system during boot.

There is also the concept of *capped* and *uncapped* partitions. A capped partition is not allowed to exceed its entitlement, while an uncapped partition is. In fact, it may exceed its maximum processor capacity. An uncapped partition is only limited in its ability to consume cycles by the lack of online virtual processors and its variable capacity weight attribute.

The variable capacity weight attribute is a number between 0–255 that represents the relative share of extra capacity that the partition is eligible to receive. This parameter applies only to uncapped partitions. A partition's share is computed by dividing their variable capacity weight by the sum of the variable capacity weights for all uncapped partitions. Therefore, a value of 0 may be used to prevent a partition from receiving extra capacity. This is sometimes referred to as a “soft cap.”

There is overhead associated with the maintenance of online virtual processors, so clients should carefully consider their capacity requirements before choosing values for these attributes. In general, the value of the minimum, desired, and maximum virtual processor at-

---

*A capped partition is not allowed to exceed its entitlement, while an uncapped partition is.*

---

tributes should parallel those of the minimum, desired, and maximum capacity attributes in some fashion. A special allowance should be made for uncapped partitions, since they are allowed to consume more than their entitlement.

In general, operating systems and applications running in shared partitions need not be aware that they are sharing processors. However, overall system performance can be significantly improved by minor operating system changes. AIX 5L V5.3 provides support for optimizing overall system performance of shared processor partitions.

The following features are available for micro-partitions:

- Dynamic LPAR extensions.
- Dynamic vary on/off of whole I/O drawers.
- Dynamic vary up/down of entitled shared processor capacity.
- Dynamic change between capped/uncapped processing.
- Dynamic vary on/off of virtual processors.
- Dynamic vary on/off (and redefinition) of virtual I/O devices.

## 7. SHARING RESOURCES BETWEEN PARTITIONS

Sharing resources between partitions is not really a new thing for IBM. It has been using this technology on mainframe systems for quite some time. Resources are shared by using a concept called virtualization, which enables physical processors and I/O devices to be shared by multiple partitions. This allows for more granular resource allocations than is possible otherwise, which leads to more partitions and higher resource utilization.

### VIRTUALIZATION OF RESOURCES

The virtualization of I/O devices is known as Virtual I/O. It allows a physical adapter on one partition to be shared by one or more partitions, enabling clients to consolidate and potentially minimize the number of physical adapters. Virtual I/O may be used to reduce costs by eliminating the requirement that each logical partition have a dedicated network adapter.

Virtual I/O is intended as a complement to dedicated or Local I/O. Partitions that have high bandwidth, low latency and high performance requirements could use Local I/O adapters to these devices. A partition can have any combination of Local and Virtual I/O. Virtual I/O support includes virtual SCSI and shared Ethernet adapters. Virtual SCSI can be used for both internal and SAN-attached disks.

Virtual I/O may be used in either dedicated or shared partitions.

The virtualization of processors, on the other hand, enables one to create a partitioning model which is fundamentally different from the POWER4™ systems, where

---

*A partition can have any combination of Local and Virtual I/O. Virtual I/O support includes virtual SCSI and shared Ethernet adapters.*

---

whole processors are assigned to partitions and are owned by them. In the new model, physical processors are abstracted into virtual processors, which are then assigned to partitions—but the underlying physical processors are shared by these partitions.

Virtual processor abstraction is implemented in the hardware and microcode. From an operating system perspective, a virtual processor is indistinguishable from a physical processor. The key benefit of implementing partitioning in the hardware is to allow any operating system to run on POWER5 technology with little or no changes. Optionally, for optimal performance, the operating system can be enhanced to exploit shared processor pools more in-depth. (For instance, by voluntarily relinquishing CPU cycles to the hardware when they are not needed.) AIX 5L Version 5.3 is the first version of AIX 5L that includes such enhancements.

## PROCESSOR SHARING

With processor sharing, more partitions can be created on a given platform, enabling clients to maximize the number of workloads that can be supported simultaneously. While it is possible to mix workloads on a single operating system image, most clients would prefer to minimize workload interactions for performance and stability reasons. Their preferred method would be to place them in separate partitions. Shared processor technology makes this easier by supporting more partitions than traditional logical or physical partitioning systems.

Another important aspect of this technology is that it results in increased physical processor utilization levels. There are two reasons for this. First, capacity can be allocated more precisely at the partition level, since

fractional allocations are supported. Second, operating systems are provided with the ability to give back capacity when there is no real work to be performed, which enables the processor to be applied elsewhere. This has the effect of minimizing idle time at the platform level, resulting in higher utilization of physical processors.

From a cost perspective, this is very significant, because it means that less processor capacity has to be purchased when consolidating servers. It is not uncommon to see 10% utilization levels, so the cost of extra processor capacity should be an important factor in the purchasing decision—particularly since processors are relatively expensive. Shared processor technology is a cost-effective alternative to server farms.

Areas to potentially benefit from inter-partition resource sharing include the following:

**Server consolidation.** For an environment consisting of a number of small servers, Micro-Partitioning is particularly attractive. A fraction of a single POWER5 processor can provide the equivalent processing capacity of one of the existing servers. Therefore, these servers can be consolidated on a single POWER5 server.

**Virtual blade servers.** Micro-Partitioning provides the ability to define hundreds of “low-cost” individual virtual blade partitions to mimic the Intel® blade server environment. In a blade environment, individual blades must have enough capacity to handle bursts of activity (Web hits), but most blades, in general, are grossly underutilized. Micro-Partitioning is a natural fit in this environment, because the idle time of a virtual blade can be utilized by another virtual blade partition.

---

*Shared processor technology is a cost effective alternative to server farms.*

---

**Production and test systems.** Micro-Partitioning provides the optimal environment for the co-existence of production and test systems. Production partitions can be defined with fixed performance requirements, such that they receive the processor capacity they require on demand. Test partitions can be defined with minimal resource commitment, but with the ability to soak up spare cycles.

**Overlapping production systems.** This is an environment where system performance is critical, but the workloads of different servers are such that the peaks in demand from one server overlap the valley of demands from another. To some degree, this environment can be serviced with dynamic LPAR. However, Micro-Partitioning provides a more granular capability which is more instantaneous.

## 8. LEVERAGING VIRTUALIZATION OF I/O AND NETWORK RESOURCES WITH AIX 5L V5.3

There are three main areas of Virtual I/O functionality available with POWER5 systems:

You need to introduce the concept of the hosting partition for Virtual I/O. You should make virtual LAN adapter and virtual disk a subset of virtual I/O. Also, Fibre Channel-attached disk subsystems (SANs) can be shared.

- Virtual Ethernet
- Shared Ethernet Adapter
- Virtual SCSI Server

Virtual Ethernet allows the administrator to define in-memory connections between partitions that appear to be high-bandwidth Ethernet connections. A virtual switch provides support for multiple (up to 4K) IEEE VLANs. Each partition can connect to multiple VLANs through one or more adapters.

Shared Ethernet Adapter is a new service that acts as a layer 2-network switch to securely route network traffic from a Virtual Ethernet to a real network adapter. It allows partitions to communicate outside the system, without having to dedicate an I/O slot and physical network adapter to that partition.

Virtual SCSI Server is a new service to serve the storage needs of client partitions. Real storage devices are attached to a dedicated partition that is called the “Virtual I/O Server partition.” The Virtual I/O Server partition defines logical volumes that are exported to client systems. On the client systems, these logical volumes appear to

---

*Virtual SCSI Server is a new service to serve the storage needs of client partitions.*

---

be SCSI disks. Virtual SCSI can be used for both internal and SAN-attached disks.

Virtual I/O may be used in either dedicated or shared partitions.

Virtual I/O slots are permanently associated with their respective partitions. Unlike physical I/O slots, virtual I/O slots are not moveable between partitions. Each slot can be populated with a virtual adapter instance (network, SCSI, or serial) and the content can be changed (redefined) without requiring a reboot. The number of slots is configurable.

Virtualization of I/O facilitates the sharing of physical resources (I/O slots, adapters and devices) between logical partitions (LPARs) on POWER5 systems. Sharing physical resources between partitions provides several key benefits.

- Allows consolidation and minimization of the number of physical adapters required.
- Provides efficient and economical I/O model, eliminating the requirement that each LPAR have dedicated physical I/O resources, such as a network adapter.
- Simplifies and drives down systems and administrative expenses.
- Enables greater number of LPARs. Typical slot requirements per LPAR are 1 slot for storage HBA, 1 slot for NIC. eServer p5 architecture is designed to support up to 254 LPARs but only 160 I/O slots. Limited slot support would limit the number of LPARs; sharing physical resources alleviates this restriction.

## 9. USING THE NEW ADVANCED ACCOUNTING FEATURE OF AIX 5L V5.3

A new Advanced Accounting subsystem is being introduced in AIX 5L V5.3. This new subsystem is inspired by mainframe technology and features many advanced capabilities, including selectable accounting records, interval accounting, data aggregation, dynamic classification of accounting data, and APIs for measuring utilization and providing accounting data.

Advanced Accounting is flexible and can be customized to meet the requirements of different computing environments. System administrators can configure the advanced accounting subsystem to produce only the records that are required by the billing solution, so that the system does not spend resources computing and reporting statistics that are not required. Selectable accounting records give system administrators greater control over the execution of their system.

### USAGE-BASED ACCOUNTING

Advanced Accounting provides usage-based information for a wide variety of system resources, so that comprehensive charge-back strategies can be developed. Accounting data can be collected for resources such as disks, network interfaces, virtual devices, file systems, memory and processors, and for work units, like processes and application-defined transactions. Not only does Advanced Accounting provide a richer source of accounting information, but in many cases, it provides better information. For example, the process record has been improved to provide an indication of the use of network resources, and CPU accounting is more accurate.

---

*Accounting data can be collected for resources such as disks, network interfaces, virtual devices, file systems, memory and processors, and for work units, like processes and application-defined transactions.*

---

Interval accounting enables accounting data to be collected at system administrator-defined intervals, so that chronological views of accounting data can be developed. This has several potential applications beyond chargeback, including capacity planning and performance analysis.

Accounting is well-suited for these activities, as it provides a comprehensive record that can be replayed at a later date, to discover what happened in the past.

From a chargeback perspective, interval accounting can be used to produce accounting bills for a partition as the billable entity, which is a very common business model. Many clients will exploit the partitioning capabilities of their POWER5 systems by creating separate partitions for each application. In this way, they will avoid performance and stability issues related to the interaction of applications. Logical partitions provide an environment that is easier to monitor, control, and measure. Advanced Accounting can be used to produce accounting records that describe the partition level use of system resources, like processors, memory, disks and network interfaces. Interval accounting is the vehicle by which this information is collected.

Interval accounting may also be used to produce intermediate process records, enabling the billing solution to produce more accurate bills, because it enables the resources consumed by the set of active processes to be factored into the bill. With the standard UNIX Accounting Utilities, it is not possible to collect information for active processes, so the business must wait for a process to exit, before it can determine the cost of a process and charge for it. This introduces a delay into the billing system, which can be considerable in some cases. Some

jobs run for months, and others never exit. Delay is not good for revenue recognition.

## DATA AGGREGATION AND DYNAMIC CLASSIFICATION

System administrators can direct the Advanced Accounting subsystem to aggregate accounting data, so that it produces less data, resulting in smaller files. Aggregation simplifies the data management aspects of post-processing and archiving accounting data. Aggregation adds like accounting records so that fewer accounting records are written to the accounting file. Aggregation is supported for records produced by vendors, processes, and application transactions. It is transparent to applications and middleware.

Advanced accounting supports the dynamic classification of accounting data, enabling the system to produce accounting data with imbedded, billable accounting information. The system administrator defines the set of billable entities, called projects, that are relevant for his business, and an accounting policy that is used to map work units, like processes to projects. These project definitions and accounting policies are then used by the system to tag processes, applications and workloads with project codes, so the accounting subsystem can automatically produce accounting records with project codes. Integration of billing information into the accounting record simplifies the post-processing phase of producing a bill from raw accounting data, enabling companies to more easily implement new billing systems—a strategic business advantage in an ever-changing business world.

Advanced accounting also provides APIs, so that the middleware and applications can measure resource uti-

---

*Advanced accounting supports the dynamic classification of accounting data, enabling the system to produce accounting data with imbedded, billable accounting information.*

---

lization and provide accounting data for their work units. Applications and middleware may use these APIs to describe the transactional nature of their workloads, so that these transactions may be factored into client bills, based on the cost of performing the transaction. Transactions are not typically known to the operating system, so middleware and applications have to be instrumented to provide this data. Advanced Accounting supports the production of accounting records from the application transaction interfaces defined by the Open Group in the document titled *Application Response Measurement (ARM) Issue 4.0 – C Binding*.

## 10. IMPLEMENTING POLICY-BASED AUTOMATIC DOWNLOADS OF UPDATES WITH SERVICE UPDATE MANAGEMENT ASSISTANT

The Service Update Management Assistant (SUMA) helps move system administrators away from the task of manually retrieving maintenance updates from the Web. SUMA offers flexible options that let clients set up policies to automate the download of fixes to their systems. SUMA policies can be scheduled to periodically check the availability of specific new fixes (APAR, PTF, or fileset), critical or security fixes, or an entire maintenance level. A notification e-mail can be sent detailing updates that are needed when comparing available fixes to installed software, a fix repository, or a maintenance level. The automation provided by SUMA will allow the time spent on these systems administrative tasks to be significantly reduced.

The SUMA command, as well as the System Management Interface Tool (SMIT) can be used to perform SUMA operations. Either a preview or download action may be specified. A SUMA policy can be set up to retrieve the following type of fixes:

- APAR ..... Specifies a request to download an APAR.
- PTF ..... Specifies a request to download a PTF.
- ML ..... Specifies a request to download a specific maintenance level.
- Fileset..... Specifies a request to download a specific fileset.
- Critical ..... Specifies a request to download the latest critical fixes.
- Security ..... Specifies a request to download the latest security fixes.
- IOserver ..... Specifies a request to download the latest IO server fixes.
- Latest..... Specifies a request to download all of the latest fixes.

## SAMPLE SCENARIOS

To create and schedule (-s) a policy that downloads the latest critical fixes on the 15th day of every month at 2:30 a.m. (using cron format), and to add a policy label through the DisplayName field (useful when listing policies through SMIT), type:

```
suma -s "30 2 15 * *" -a RqType=Critical -a DisplayName="Critical fixes - 15th Monthly"
```

To run a task immediately (-x) that downloads the entire 5300-02 maintenance level into the /lppsrc/5302 directory, while turning off filtering against installed software (FilterSysFile field), and running an lppmgr clean operation after the download to remove any superseded updates, duplicates base levels, and conflicting updates, type:

```
suma -x -a Action=Clean -a RqType=ML -a RqName=5300-02 \
-a DLTarget=/lppsrc/5302 -a FilterSysFile=/dev/null
```

Highlights of this new feature include:

- Moves administrators away from the task of manually retrieving maintenance updates from the Web.
- Provides clients with flexible options.
- Policy can be scheduled to run periodically. (For example, download the latest critical fixes weekly.)
- Can compare fixes needed against software inventory, fix repository, or a maintenance level
- Receive e-mail notification after a fileset preview or download operation.
- Allows for ftp, http, or https transfers.
- Provides same requisite checking as the IBM fix distribution Web site.
- Available through SMIT menus (smitty suma) or a command line interface.

## 11. RECOMPILING LINUX APPLICATIONS TO RUN IN A NATIVE AIX 5L ENVIRONMENT

There is a high degree of Linux application affinity provided with AIX 5L. This means that a wide variety of applications built for Linux can be rebuilt to run natively on AIX 5L with minimal or sometimes no porting effort. This is possible predominantly for two reasons. First, AIX 5L provides UNIX industry-standard application programming interfaces (API's), supplemented with additional Linux OS-specific API's added to AIX 5L, specifically to increase application portability between the operating systems. Second, a collection of common Linux development utilities is available to install and run natively on AIX 5L. This collection, the AIX® Toolbox for Linux Applications, is provided on a separate companion CD with each shipment of AIX 5L. These utilities are also available as free downloads from an IBM FTP site. See the Web page, <http://www.ibm.com/servers/aix/products/aixos/linux/> for more information.

---

*A collection of common Linux development utilities is available to install and run natively on AIX 5L.*

---

Of course, there is such a broad assortment of software applications available on Linux that a single description cannot capture all of the different experiences of porting an application from Linux to AIX 5L. Still, it is possible to make some generalizations that cover the most common cases. When building an application on Linux, a typical case is first to run the configure script provided with the application, which determines availability of assorted commands and interfaces on the system. This is normally followed by running the “make” command, which on Linux is simply the GNU version of the make utility. (GNU is a collection of free software utilities managed by the Free Software Foundation; most of the basic commands on a Linux system are GNU utilities.) During this build

phase, the source is compiled using the GNU compiler, “gcc.” Finally, when the build has completed, the application is often packaged into an installable image called an RPM (the “RPM Package Manager”). This image can then be installed on a Linux system.

On AIX 5L, because of the Linux affinity features, the process is almost identical. Some applications may simply build just as they are; however, for most, you will likely need to install at least some of the utilities from the aforementioned *AIX Toolbox* onto the AIX 5L system. For example, the “make” utility is available on the *AIX Toolbox*. It is true that AIX 5L provides its own native version of “make”, but the *AIX Toolbox* provides the same GNU “make” that is found on Linux systems. The native “make” might work fine with a given application, but using the GNU “make” ensures that any subtle behavior differences are avoided and that the command’s behavior matches what is encountered on Linux. You can enforce the use of the GNU “make” on AIX by changing your environment PATH to include /usr/linux/bin before /usr/bin or /bin. This model extends to the other utilities in the *AIX Toolbox*. For example, by installing packages like “coreutils”, “autoconf”, “automake”, “libtool”, “install”, etc., you will have available the same commands with the same behavior, that are, found on Linux. Once you have installed the development tools from the *AIX Toolbox* and updated your PATH variable, you will be ready to build your application on AIX 5L, using the same procedure that you used on Linux. Even the RPM utility is available on AIX 5L, preinstalled, so you can create packages in RPM format and install them in exactly the same way that you do on Linux.

---

*Even the RPM utility is available on AIX 5L, preinstalled, so you can create packages in RPM format and install them in exactly the same way that you do on Linux.*

---

In some cases, you might find that you have an application which depends on a Linux system library, which is

not available on AIX 5L. For example, your application might depend on the zlib library, or the db library. While these libraries (and many others) are not provided natively on AIX 5L, they are included in the *AIX Toolbox*. Just like the utilities, you can install any such libraries that your application requires, and then proceed with your build, following the normal build routine. In some rare cases, it is possible that you might find a library or utility that you need which does not appear in the *AIX Toolbox*. If so, the same procedures can be used to build your own copy of the missing piece, just as if you were porting another application.

Finally, a note on compilers: While there is a version of “gcc” available in the *AIX Toolbox*, it is recommended that you attempt to compile your application using the native VisualAge IBM compilers, if you have them. Normally, this transition will work smoothly, though in some cases, minor code patches may be needed to account for small behavior differences. Because VisualAge is written and optimized for AIX 5L specifically, it generally produces smaller and faster executables than with “gcc”. If performance is not an issue, however, “gcc” may be used to most closely match the Linux build experience. Also note that there are many versions of “gcc” around. While the *AIX Toolbox* provides one version, you may prefer a different one. If you do an Internet search you should be able to find other versions of “gcc”, or you could even download the source from GNU and build a copy.

The descriptions above apply to a majority of Linux applications that an end user might want to run. However, there are always exceptions. Packages such as Linux device drivers or kernel extensions are very specific to the operating system kernel itself, and generally require some degree of code change—as would be expected in ports

between any two different operating systems. Of course, the vast majority of application code does not fall into this category.

## 12. ADDITIONAL RESOURCES

### IBM AIX 5L INFORMATION

AIX 5L is rapidly emerging as the preferred platform for UNIX users and independent software vendors. Learn more about AIX. Visit <http://www.ibm.com/servers/aix/>

### AIX WHITEPAPERS

Download your choice of numerous complimentary AIX whitepapers. Titles include: Advanced Accounting for AIX 5L Version 5 and IBM eServer p5 - AIX 5L Support for Micro-Partitioning and Simultaneous Multi-threading. Updated frequently.

Visit [http://www-1.ibm.com/servers/eserver/pseries/library/wp\\_aix\\_lit.html](http://www-1.ibm.com/servers/eserver/pseries/library/wp_aix_lit.html)

### LOGICAL PARTITIONING (LPAR) WHITEPAPERS

Download a wide range of complimentary whitepapers specifically related to LPAR and IBM eServer pSeries systems. Titles include: LPAR for Decision Makers and LPAR Heterogeneous Workloads on the IBM pSeries 690 System.

Visit <http://www-1.ibm.com/servers/eserver/pseries/lpar/wp.html>

### REDBOOKS

IBM Redbooks are developed and published by IBM's International Technical Support Organization (ITSO). They are designed to deliver technical know-how to technical professionals. Download your choice of these complimentary Redbooks in PDF format:

AIX 5L Performance Tools Handbook

<http://www.redbooks.ibm.com/abstracts/sg246039.html>

CSM Guide for the PSSP System Administrator

<http://www.redbooks.ibm.com/abstracts/sg246953.html>

Developing and Porting C and C++ Applications on AIX

<http://www.redbooks.ibm.com/abstracts/sg245674.html>

AIX 5L Differences Guide Version 5.2 Edition

<http://www.redbooks.ibm.com/abstracts/sg245765.html>

IBM eServer Certification Study Guide - AIX 5L Installation and System Recovery

<http://www.redbooks.ibm.com/abstracts/sg246183.html>

IBM eServer Certification Study Guide - AIX 5L Performance and System Tuning

<http://www.redbooks.ibm.com/abstracts/sg246184.html>

IBM eServer Certification Study Guide - AIX 5L Problem Determination  
Tools and Techniques

<http://www.redbooks.ibm.com/abstracts/sg246185.html>

IBM eServer Certification Study Guide - AIX 5L Communications

<http://www.redbooks.ibm.com/abstracts/sg246186.html>

AIX Reference for Sun Solaris Administrators

<http://www.redbooks.ibm.com/abstracts/sg246584.html>

A Practical Guide for Resource Monitoring and Control (RMC)

<http://www.redbooks.ibm.com/abstracts/sg246615.html>

Introducing VERITAS Foundation Suite for AIX

<http://www.redbooks.ibm.com/abstracts/sg246619.html>

AIX and Linux Interoperability

<http://www.redbooks.ibm.com/abstracts/sg246622.html>

AIX Version 4.3 to 5L Migration Guide

<http://www.redbooks.ibm.com/abstracts/sg246924.html>



The IBM logo, the e-business logo, eServer, POWER, pSeries, AIX 5L and Micro-Partitioning are trademarks or registered trademarks of IBM Corporations and its subsidiaries. UNIX is a registered trademark of The Open Group in the United States and other countries. Other company, product and service names may be trademarks or service marks of others. Linux is a registered trademark of Linus Torvalds in the United States and other countries. Other company, product and service names may be trademarks or service marks of others. © IBM Corp. 2004