Table of contents

IBM® eServer™ Solutions Development

IBM POWER5 and eServer Trends and Directions

Introduction

The IBM® POWER™ processor architecture is much more than just a family of chips. It is a comprehensive systems-design approach derived from a single architectural framework. The POWER processor provides a systems-based control point that is structured to integrate technology across the solution for maximum success as a unified architecture. As a result, the POWER product base is diverse, ranging from game consoles and blade servers to desktops and supercomputers. In this online course, we will provide an overview of the IBM POWER trends and directions, with primary focus on the capabilities and advantages of the IBM POWER5™ processor in the IBM eServer™ family of products.

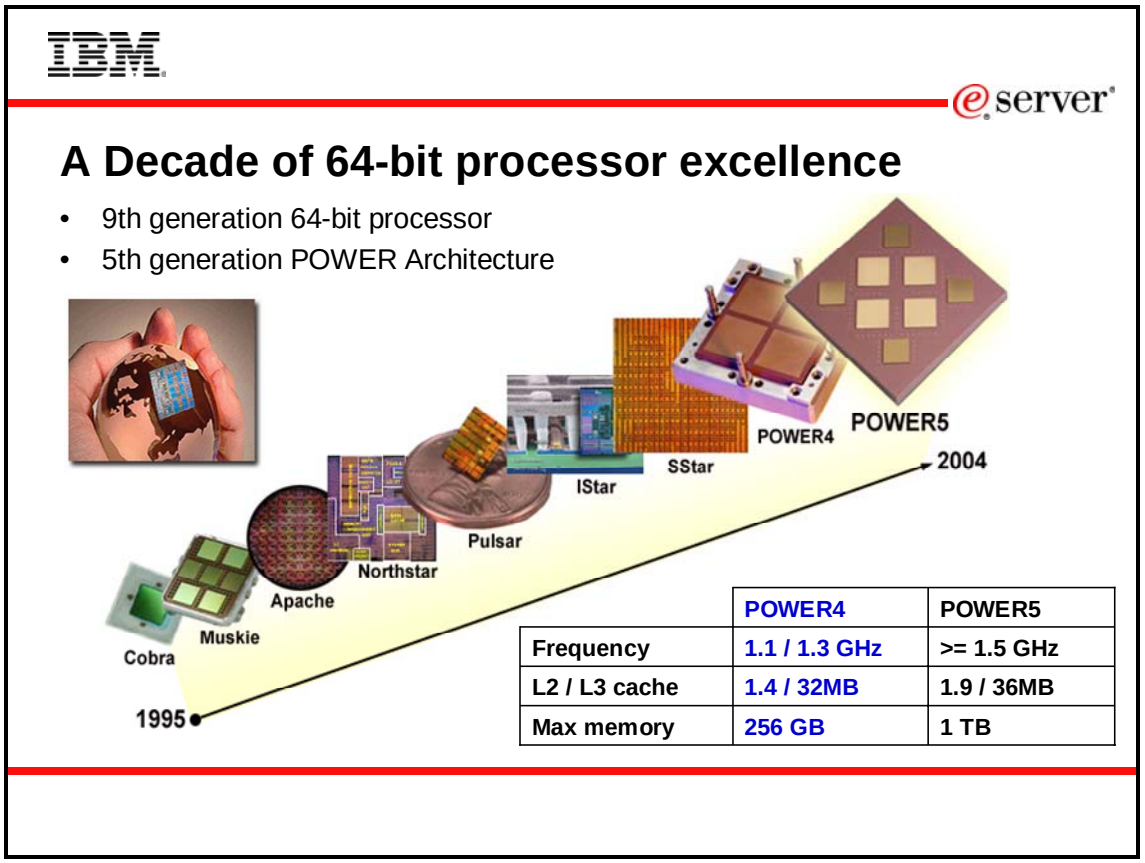

A Decade of 64-bit processor excellence

The POWER5 chip is the ninth generation of 64-bit processors delivered on IBM servers since 1995, and as its name suggests, is also the fifth generation of the POWER processor based on the IBM Power Architecture™. The first generations of the POWER processor drove IBM RS/6000® servers. After this, a workstation was optimized for scientific and compute-intensive performance through POWER. Later, the processor designs used in the RS/6000 and IBM AS/400® servers were gradually converged, bringing together commercial and compute-intensive processing characteristics.

Starting with the AS/400 systems delivered in 1997, the same processors were used on both platforms, further converging the development investments between the RS/6000 and AS/400 product lines. Now with POWER5, IBM delivers a truly common platform between the two servers, which have become the IBM eServer p5 and IBM eServer i5 server families, respectively.

Major enhancements have been introduced since IBM combined all of the company's 64-bit processor development efforts into the Power Architecture. Larger caches on processor chips, fast and wide memory bus interfaces, and new caching algorithms (which optimize processor utilization by bringing instructions and data faster into the processors) resulted in superior performance in commercial and noncommercial applications.

The latest advances of the Power Architecture implemented with the POWER5 processor include simultaneous multithreading (SMT), which allows the execution of two

threads on a single physical processor at the same time. This provides a major boost in performance. Multiple physical processors are combined on a single processor chip module to allow much faster communication between processors and to leverage information stored in L2 caches across processors, resulting again in dramatic performance improvements.

IBM is an experienced provider of tried-and-true 64-bit processor technology. The next versions of the Power Architecture-based processors are already on the drawing board and will provide even more advanced features and functions for the IBM eServer product family.

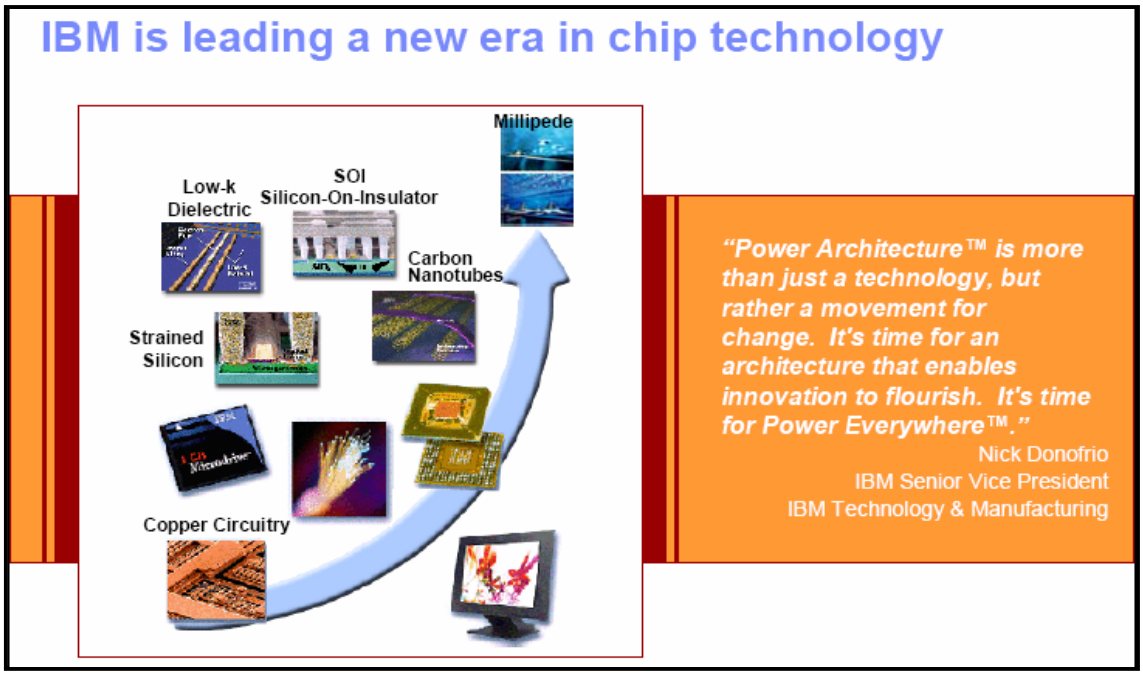

IBM is leading a new era in chip technology

Over the years, chips and packaging technology from the IBM Microelectronics Division have contributed to a significant number of IBM patents and have helped position IBM as a technology leader. To further strengthen and support our chip technology leadership position, IBM opened the world's most advanced semiconductor facility in East Fishkill, New York, in 2002. For the first time anywhere, this new facility combines IBM chip-making breakthroughs, such as copper interconnects and silicon-on-insulator (SOI) technologies. Importantly, this plant also supports low-k dielectric insulation on 300-millimeter (12-inch) wafers, thus allowing for the creation of chips with circuits smaller than 100 nanometers. These chips contain about 50 billion transistors. This $2.5 billion plant represents the largest capital investment in IBM history. One key result of this investment in semiconductor technology is the development and delivery of the Power Architecture.

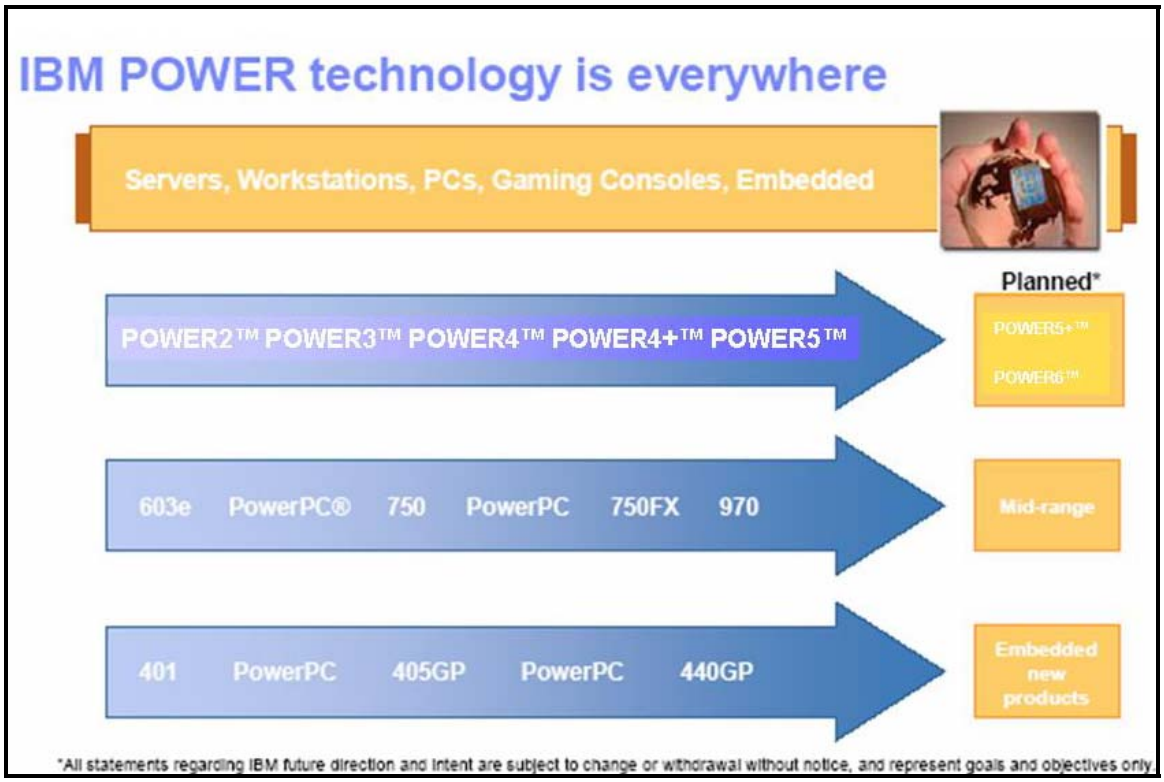

IBM POWER technology is everywhere

As you will see by looking at this screen, the IBM Power Architecture has evolved into more than just a semiconductor chip. Having been built from one framework, it delivers a systems-based control point for integrating and unifying technologies throughout the enterprise. As mentioned in the introduction to this course, the IBM Power Family™ of products has a broad reach, supporting game consoles, servers, desktops, and supercomputers. Up to now, this diversity is one of the best kept secrets for IBM, which has been steadily moving into position as an industry-leading, 64-bit architecture supplier. The Power Architecture is being deployed in high volumes in many areas of business, and is not limited to IBM servers. POWER is everywhere, including the IBM eServer platform and workstations, the IBM Blue Gene® supercomputers (including Blue Gene/L, the world's fastest supercomputer), and even in the three leading game consoles: the Microsoft® Xbox™, Nintendo® GameCube™, and Sony® PlayStation®.

(For more information on Blue Gene technology, refer to the link provided in the "Additional information" section at the end of this course.)

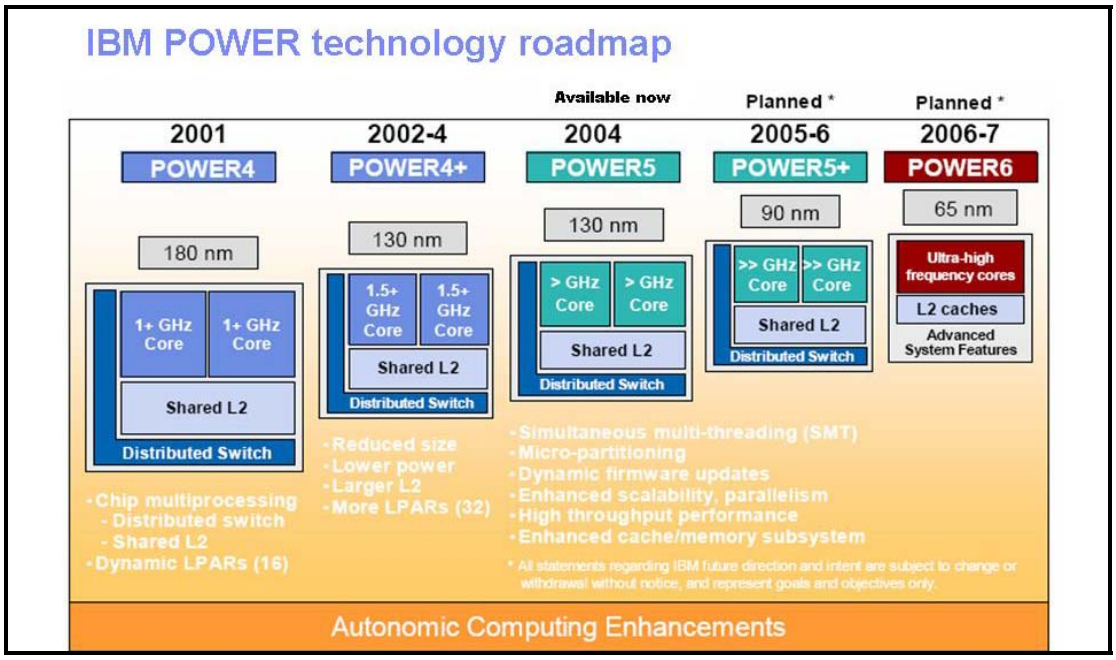

IBM POWER technology roadmap

This screen shows the POWER processor roadmap as it is planned through 2007. As you can see, the list of Power Architecture industry firsts continues to grow with each new processor release. IBM plans to ship POWER processor-based servers on time and on schedule with technology that is generations ahead of anything else in the marketplace. For solid, proven architectural superiority, POWER is, without question, one of the soundest investments an enterprise can make for present and future infrastructure success.

The Power Architecture is the foundation of the IBM eServer i5 and iSeries™, IBM eServer p5 and pSeries®, and the IBM OpenPower™ families of servers.

Look at the list of POWER5 processor enhancements shown on this screen. These enhancements will allow users of the IBM eServer line of products to achieve new levels of system utilization and productivity. These enhancements also continue to deliver on the IBM strategy that supports On Demand Business™.

For example, consider POWER5 features such as simultaneous multithreading (SMT), which allows a single processor to behave as if it were two processors running at full speed. IBM Micro-Partitioning™ technology allows up to 254 partitions on a 64-way (64-processor) server. Self-detecting and self-healing circuitry identifies data transmission errors and corrects them automatically.

POWER5 Multichip Module (MCM)

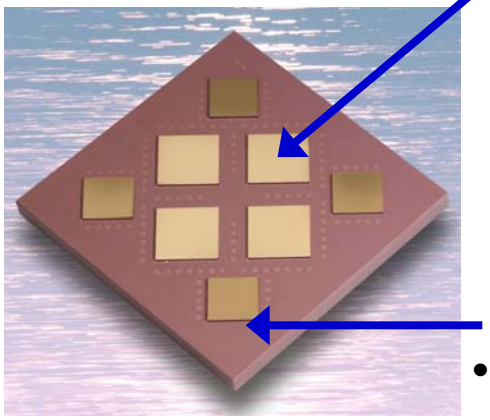

- POWER5 chip
  - —4 chips (8 cores) per MCM
  - —2 cores per chip
  - —4 logical processors per chip with simultaneous multithreading (SMT)
  - —Up to 1.9 MB shared L2 cache per chip
  - —Integrated L3 cache & memory controller
- L3 cache
  - —Up to 36 MB
  - —Up to 144 MB per MCM

POWER5 multichip module (MCM)

IBM continues to deliver more power and function into smaller and smaller footprints with each generation of POWER. With POWER5, IBM can package up to four POWER5 chips on one multichip module (MCM). Every chip consists of two processor cores, enabling a single MCM to contain up to eight processors. With SMT, each of these processors can be viewed as two logical processors. *(SMT is described in more detail on the next screen.)* Up to 1.9 megabytes of shared L2 cache and 144 megabytes of L3 cache can be packaged with each MCM.

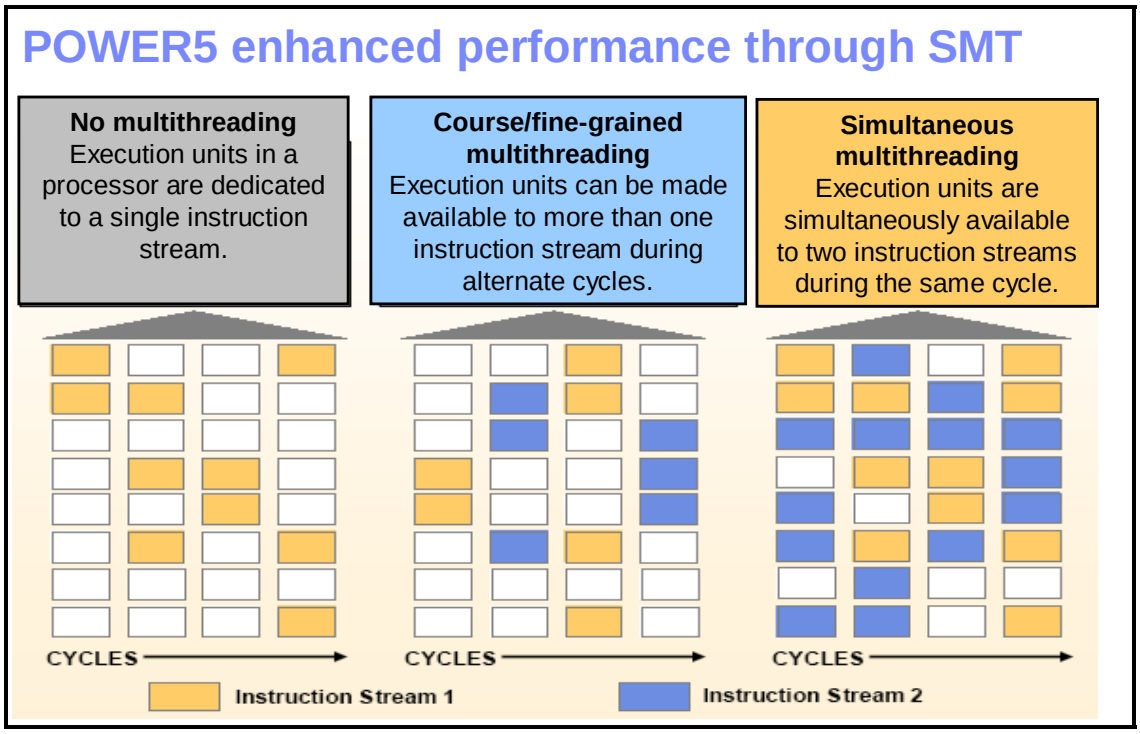

POWER5 enhanced performance through SMT

POWER5 servers from IBM offer a number of means to further enhance productivity and performance, starting with SMT, which is proven to increase system-level throughput 30 to 40%, depending on workload. (The “Additional information” section of this course includes references to other materials that discuss throughput gains related to SMT in greater detail.)

Without multithreading, only one instruction path can be processed per cycle by the execution units in a microprocessor.

With the coarse or fine-grained versions of multithreading available in earlier POWER implementations (for example, the multithreading available on RS/6000 S85 servers), it is possible for two instruction paths to alternately access the microprocessor execution units.

Now, with the advent of SMT provided by POWER5, **two instruction paths can share access to execution units on every cycle**, thus increasing system-level performance, utilization, and throughput.

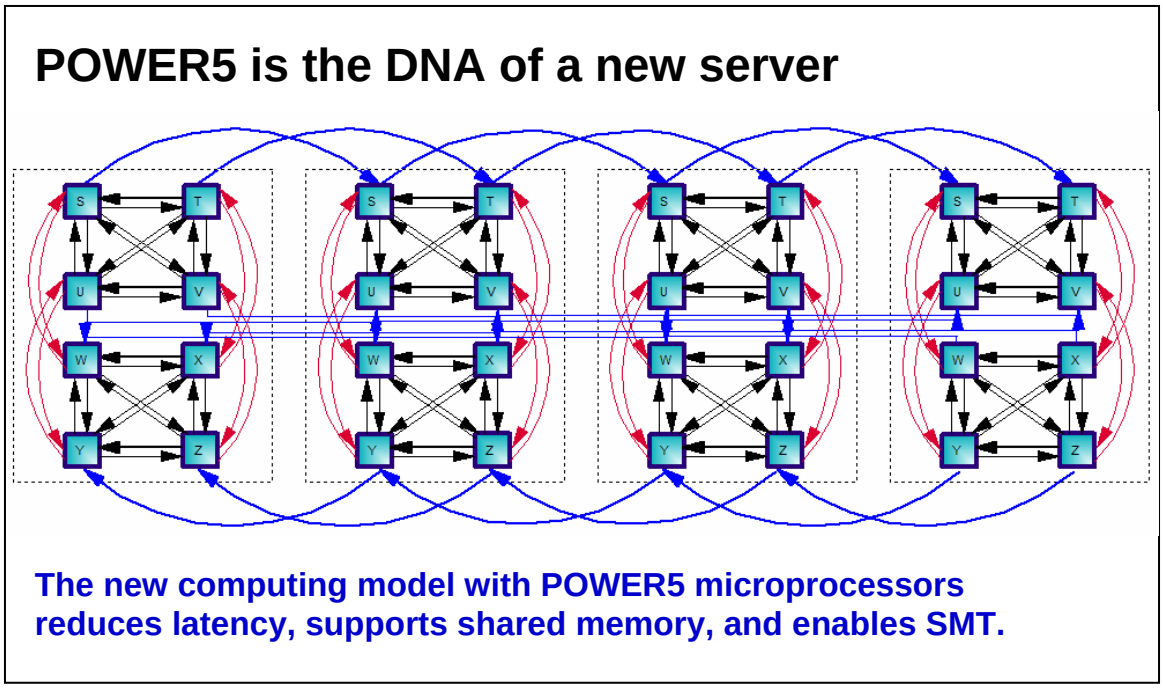

POWER5 is the DNA of a new server

What about processor topology? The largest IBM eServer POWER4™-based servers consisted of 32 processors. But POWER5 servers can now support up to 64. These processors are organized within a set of 8-processor MCMs or 2-processor Dual-Chip Modules (DCMs). Servers that are even larger than 8-processors are built using multiples of these MCMs. Within an MCM, there are four processor chips, each chip supporting two physical processors.

Both of the processors on a chip share the same Level 2 (L2) cache, but each processor has separate Level 1 (L1) instruction and data caches. A processor executes fastest when the accessed data and instruction stream are in the L1 cache. If the needed data or instruction stream is not in the L1, the processor effectively waits for a number of cycles while a block of storage is loaded into an L1 cache from the L2. The L2 cache is considerably larger than the L1 cache.

These servers also support a Level 3 (L3) cache, an equal portion of which can be associated with each processor chip and MCM. The L3s, which can feed the L2s in the event of an L2 cache miss, are considerably larger than the L2s. Similar to the performance decrease when feeding L2 cache after an L1 cache miss, L3s take considerably longer to complete an access.

The contents of the L2 and L3 caches are called store-in caches because a change to a block of storage can be held in an L2 or L3 indefinitely before finally written back out to main storage. Although by no means precise, you can assume the following:

- It takes approximately 10 to 20 cycles to fill the L1 from the processor's L2
- It takes about 100 cycles to fill the L2 from the processor's L3
- It takes a few hundred cycles to complete an access of main storage

Clearly, a program’s performance is largely dictated by the processor’s ability to hold data and instruction streams in its cache.

The contents of all SMT caches are called "coherent." A changed block of storage residing in one processor’s cache is visible to any other processor. For example, when processor A executes a read from the block of storage changed by processor B, it can get to the changed block by accessing the cache storage of processor B, even if A and B are in different MCMs. An unchanged block of storage can be held in the cache of multiple processors. However, upon a request to change that block, the block is effectively removed from all those caches and held only in the cache of the processor making the change. It is convenient to think of these blocks of storage and data packets as being rapidly pulled and copied from processor cache to processor cache.

L1 and L2 cache lines hold the contents of 128-byte blocks of main storage aligned on 128-byte boundaries.

It is this set of POWER5 design features that support server scalability and growth.

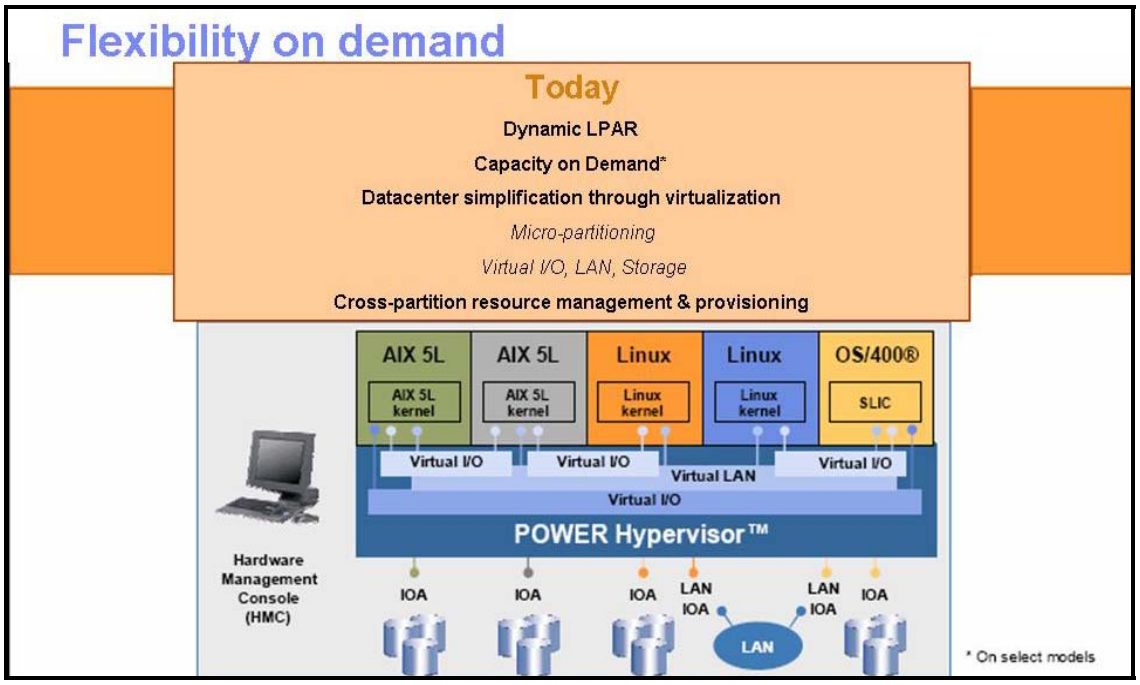

Flexibility on demand

Now that we have discussed the POWER5 processor design and packaging, let us examine how IBM incorporates and leverages the capabilities provided by these POWER5 building blocks in its IBM eServer line of products.

The ability to manage an enterprise with flexible IT solutions is increasingly important in today's dynamic marketplace. An organization must be able to adjust to unpredictable demands effectively. The IBM eServer On Demand Business strategy is fundamentally about flexibility. The strategy is to provide businesses with the capability to deliver the right resources to the right business processes at the right time.

IBM eServer systems answer the call for flexibility in many ways. With dynamic logical partitioning (LPAR), companies can immediately create new virtual servers to enable rapid application deployment. With Capacity On Demand features, resources can be dynamically increased or decreased when needed, without disrupting the business. Together, dynamic LPAR and Capacity On Demand provide the ability to allocate processing power where it is needed in a flexible and efficient manner.

The flexibility story has improved dramatically by leveraging POWER5 capabilities. The iSeries, pSeries, and OpenPower servers now enjoy much more of the same superior utilization available on mainframes. The POWER5 processor itself is a virtualization enabler that allows uncapped partitions in shared processor pools, as well as up to 10 logical partitions per processor, with a potential total of 254 partitions in a single server. Virtual I/O, LAN, and storage all complement POWER5 virtualization to provide a broad spectrum of virtual resource integration and workload flexibility.

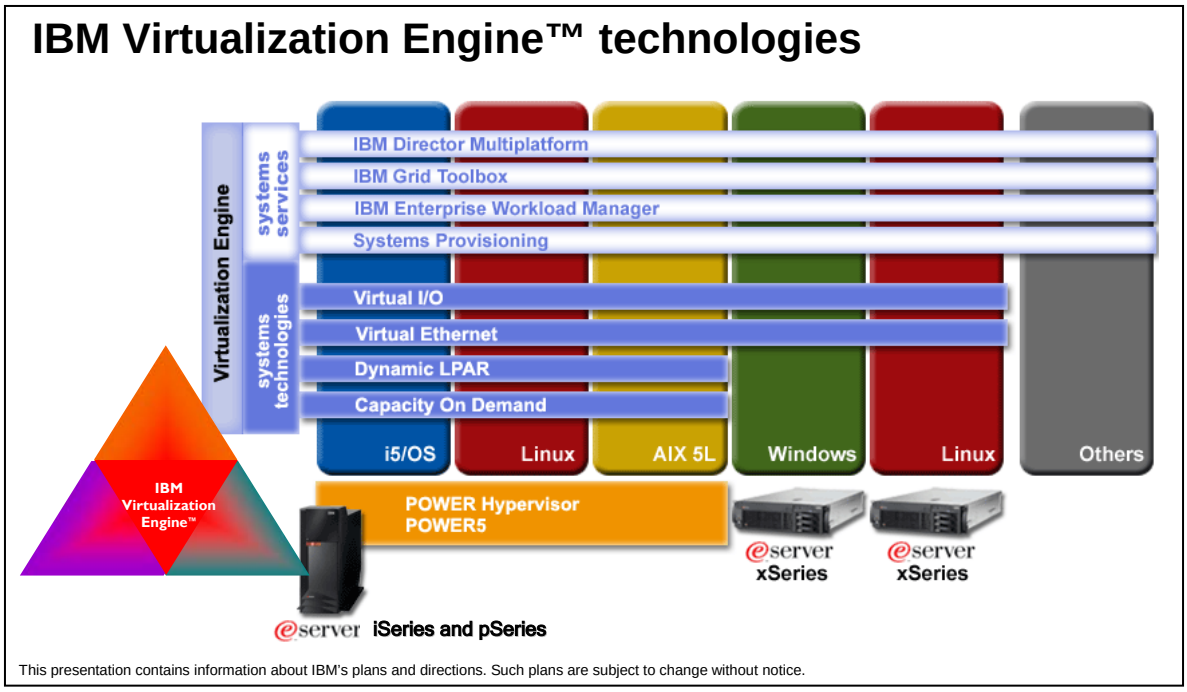

IBM Virtualization Engine technologies

Managing the on demand operating environment can be challenging. This has driven the development of the innovative IBM Virtualization Engine™ technologies and services that are aimed at redefining the economics and management of the operating environment. Virtualization Engine technologies provide a consolidated view and easy access to resources in a network, and allow multiple operating systems on a single server. With Virtualization Engine technologies, businesses will be able to extract value from billions of dollars worth of underutilized servers and storage.

IBM is previewing some of its newest technologies to simplify the management of hundreds of servers into a single view. This unified management view will enable and encourage significant increases in server utilization. By setting a standard for how servers interact, IBM provides the ability to get more work done with existing hardware. These capabilities will be built into IBM eServer and TotalStorage® servers, making them easier to set up, manage, and connect with other IBM and non-IBM servers.

Idle server boxes rack up costs and complexity. With server and storage prices continuing to drop, businesses are buying more servers. This can result in greater IT inefficiencies and complexities because companies end up hiring more people to support and manage these servers. However, distributed server utilization rates are typically extremely low, averaging only 10 to 20%. This is very inefficient and costly. Aimed at improving IT economics and simplifying operations, Virtualization Engine technologies can increase infrastructure utilization by as much as 80%.

Heterogeneous servers managed as one system has long seemed an unattainable nirvana. Virtualization Engine technologies also allow workloads to be managed across distributed computers or on a grid. If a server failure or disaster occurs, these technologies can draw from available resources, helping to minimize downtime. New Virtualization Engine components can also make one server act like many. The business

can run as many as 10 servers per microprocessor, with the potential of turning a 4-way server, for example, into a 40-way server that supports one or multiple operating system types or versions at the same time. This encourages much more efficient use of the server.

This is not a novel approach in the mainframe arena, but for distributed UNIX® environments, it is a big deal. One processor in a server can be partitioned into as many as 10 logical servers, each capable of running multiple instances of the same or different operating systems. For example, on the pSeries platform, a 16-way server can run as many as 160 server partitions. One hundred LPARs can be running the Linux™ operating system, 40 LPARs can be driven by the IBM AIX® operating system, and 20 LPARs can be executing under the IBM i5/OS™ operating system (known as OS/400® prior to Version 5 Release 3).

This represents continued efforts by IBM to implement core technologies from the mainframe down into its midmarket servers. In addition, key technologies, such as basic provisioning from IBM Tivoli®, coupled with grid innovations and base infrastructure technologies from IBM WebSphere®, have been built into the family of Virtualization Engine technologies.

Virtualization Engine technologies are available now on IBM eServer iSeries, pSeries, and OpenPower servers.

(**Note**: Refer to the “Additional information” section for a link to the eServer Software Information Center for more content on the benefits of these technologies.)

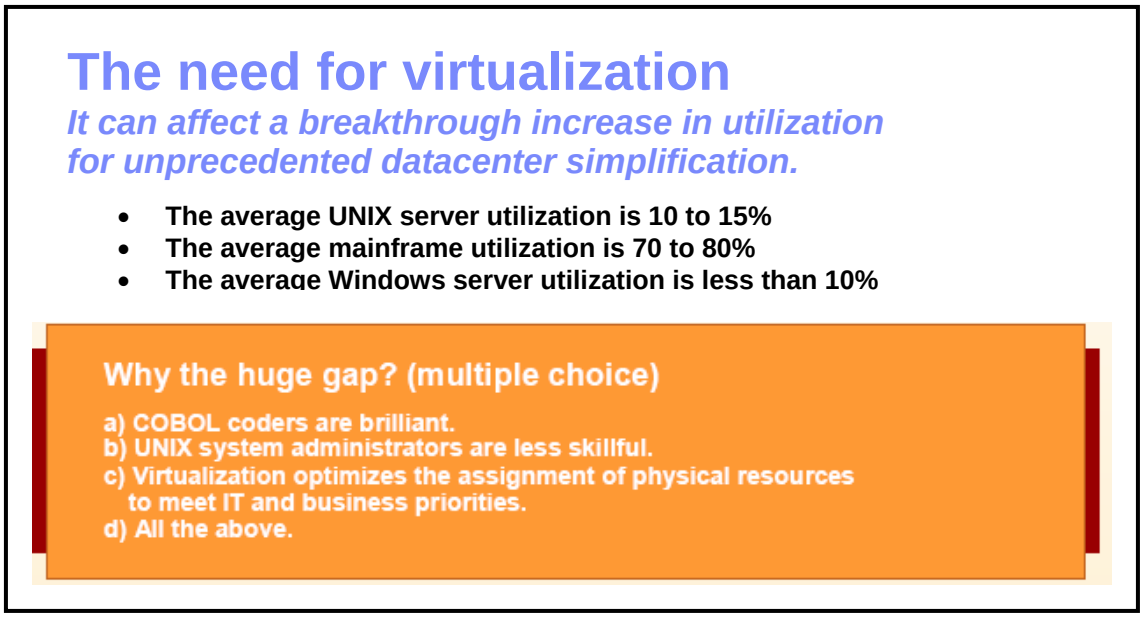

The need for virtualization

A recent study found that though the typical mainframe runs at 70 to 80% utilization, the average UNIX system is utilized to only 10 to 15% of its capacity. The picture becomes even worse for Microsoft Windows® servers, where average utilization falls under 10%.

See the multiple-choice question in the screen. Why is there such a huge gap in utilization rates between platforms? The answer is ‘C’ obviously! Virtualization, which is a key component of mainframes, optimizes the assignment of physical resources to meet IT and business priorities.

For years, larger IT shops have focused on getting the most out of their mainframes. For these companies, wasted CPU cycles meant wasted time and money. Therefore, investment and effort went into ensuring that every possible byte of capacity could be squeezed out of their systems.

Virtualization has become the new game changer. It is the technology set that is redefining how businesses implement IT solutions across the enterprise and simplify the complexity of data centers.

(**Note**: For more details on the utilization statistics shown here, see the white paper: *Scorpion — Simplifying the Corporate IT Infrastructure*. A link to this paper is contained in the “Additional information” section of this course.)

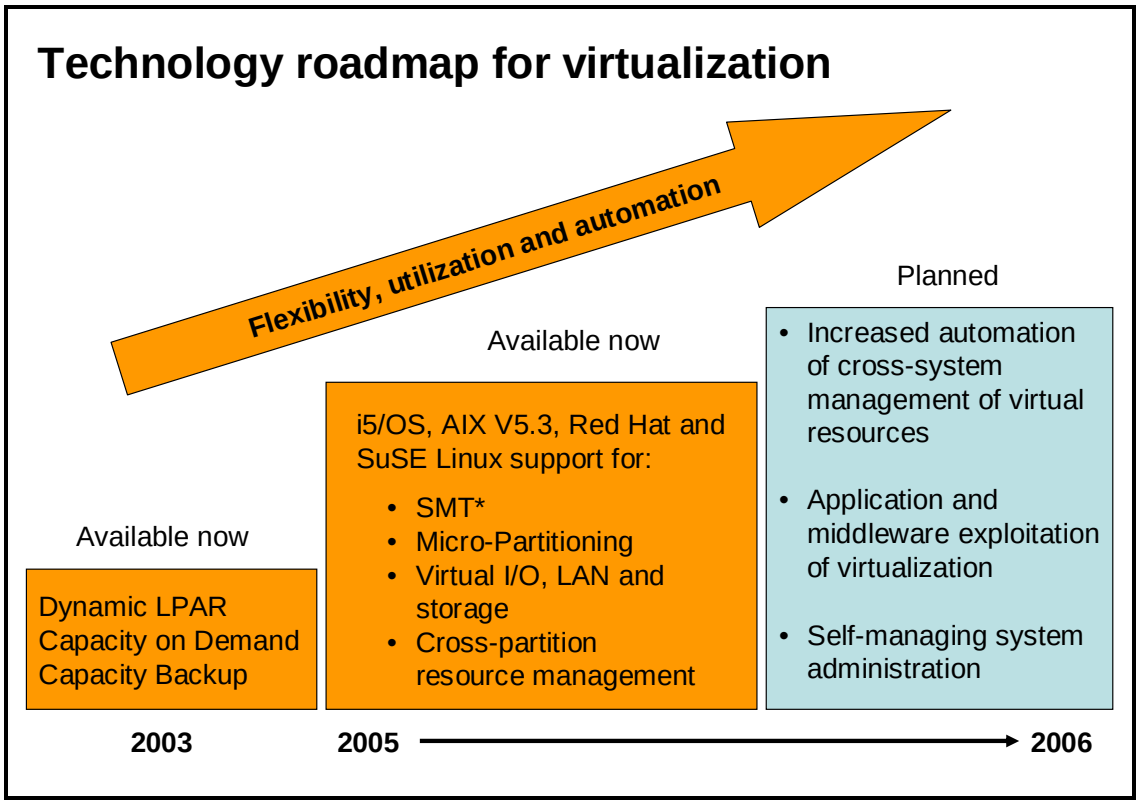

Technology roadmap for virtualization

Eventually, the on demand world will become a virtual world because the economics make sense. In a virtual world with a pay-as-you-go pricing model, underutilization is a thing of the past.

Virtualization offers the prospect of increasingly efficient server utilization and higher productivity. Virtualization enables the rapid deployment of IT resources that support business-critical applications. It does this through flexibility in the allocation of resources, higher and more cost-effective system utilization, and increasing automation of server administration tasks.

From now and until the end of the decade, IBM will continue to deliver advanced virtual capabilities that enable the IBM eServer platforms to be more and more self-managing. This began with the eServer announcements in 2004 and 2005.

Let us look more closely at some of these new capabilities in more detail.

Dynamic logical partitioning

- Includes POWER Hypervisor™ that supports i5/OS, AIX 5L™, and Linux in up to 254 partitions
- Improves server utilization rates across multiple workloads
  - Automatic processor balancing with uncapped partitions
- Increases fault tolerance and lowers partition management costs
  - Primary partition replaced by Hardware Management Console (HMC) on iSeries
- *An IBM Virtualization Engine systems technology*

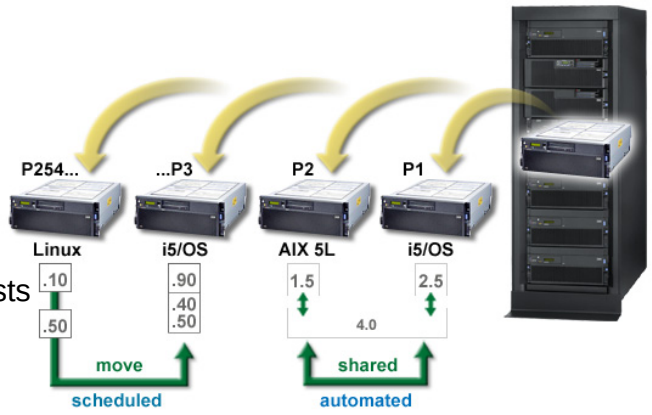

Dynamic logical partitioning

The IBM eServer platforms, combined with the new IBM POWER Hypervisor™ systems management console, feature exceptional support for dynamic logical partitioning (DLPAR). This support includes the latest enhancement, which enables automatic processor movement between uncapped partitions. IBM eServer models can support up to 254 partitions, with a maximum of 10 per processor. Also, the POWER Hypervisor can support partitions running a variety of operating systems, including AIX 5L™, i5/OS, and Linux.

Today, many high-end iSeries clients are already exploiting this mainframe-class LPAR technology on their n-way systems. LPAR usage is also increasing on single-processor IBM eServer servers.

The introduction of POWER Hypervisor and the IBM Hardware Management Console (HMC) eliminates the requirement to define a primary partition for LPAR management. This improves total server availability and reduces scheduled downtime, because all partitions can be started or restarted individually without affecting the entire system. The POWER Hypervisor can also create new partitions dynamically without requiring a server restart.

The POWER Hypervisor also features the capability to create a new kind of partition, an uncapped partition, in addition to dedicated partitions and capped partitions. Uncapped partitions enable businesses to maximize server utilization rates by automatically moving any unused processing resources, defined in a shared processing pool, to other partitions.

For example, the following four partitions are defined as:

- Partition 1 has two dedicated CPUs
- Partition 2 has .75 processing units defined as capped
- Partition 3 has .50 processing units defined as uncapped
- Partition 4 has .75 processing units defined as uncapped
    - Partition 4 also has the highest priority ranking for consuming any unused processor capacity.

In this scenario, when Partition 4 reaches 100% utilization of its allocated CPU, the POWER Hypervisor will look for more resources in partitions 2 and 3. If there are any unused processing units, the POWER Hypervisor will automatically move them into Partition 4.

One potential difference between capped and uncapped partitions pertains to software licensing. Suppose you wanted to run an application (Application A) in a shared processor pool (Pool B) that contains four processors. Assume, too, that Application A would never need more than one processor. You can create a special capped partition for Application A, which gives you the granularity to allocate as little as one tenth of a processor or as much as an entire single processor. Any unused capacity in this partition will be used by the shared pool and any other partitions.

Why create a single processor capped partition for Application A in this example? Because the upper processor limit in an uncapped partition is defined as the number of processors available in the entire shared pool. Hence, to place Application A into the 4-processor shared pool partition, you must purchase a 4-processor license for Application A, because it is possible that all four shared processors might be allocated to the application at the same time.

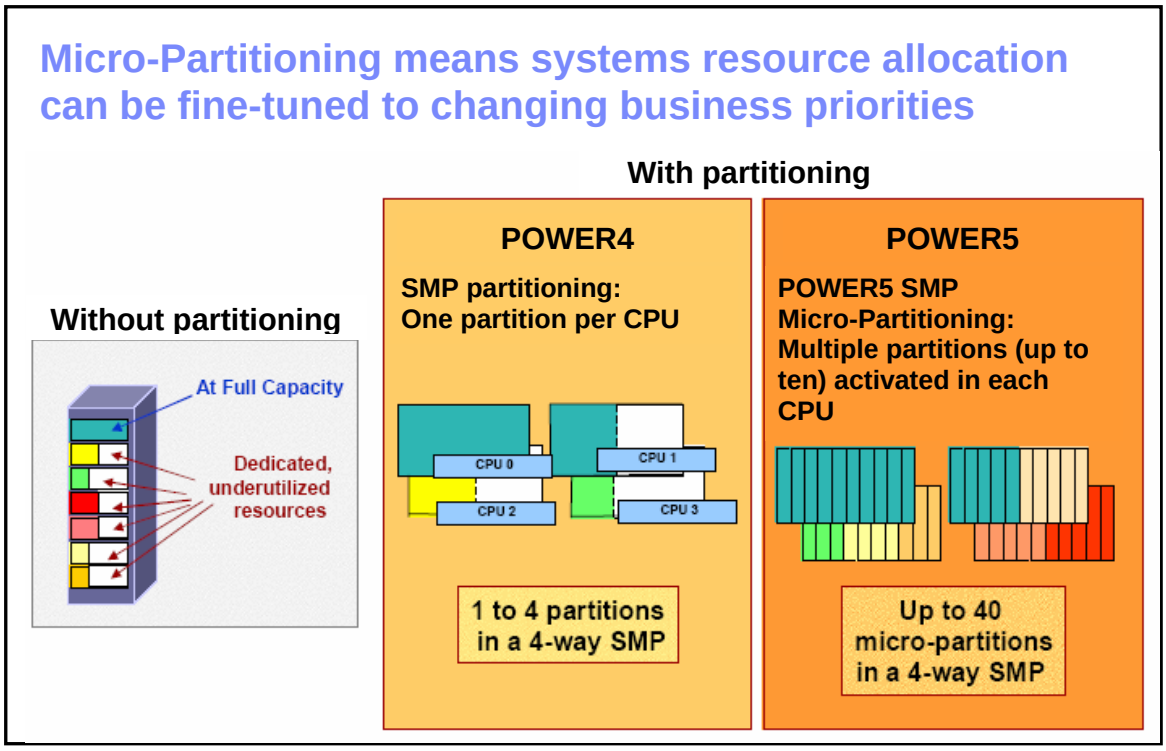

Micro-Partitioning technologies

Another major flexibility enhancement to the IBM eServer platform is provided by POWER5 technology through its Micro-Partitioning innovation, which enables the creation of multiple virtual partitions within a single processor. Each virtual partition can be tailored to the resource requirements of a particular application based on business needs and priorities. This results in higher levels of sustained utilization, better system level throughput, and greatly increased IT productivity.

Previous partitioning support required allocation of one or more entire microprocessors for each partition supported. Depending on the nature of the application supported, partition resources (processor cycles, memory, or I/O) could have been underutilized, resulting in an unnecessarily high total cost of ownership (TCO).

With the POWER5 SMP Micro-Partitioning design, you can allocate partial microprocessors to better match workload and increase utilization. Micro-partitions can be tailored to the demands of individual applications, in increments of one hundredth of a processor.

Hardware Management Console (HMC)

- Single console for POWER5 servers
  - Pre-installed Linux-based workstation
  - Supports local consoles, including 5250
  - Web-based System Manager enables local or remote management for HMC control and status
- LPAR and CoD
  - Now configured via HMC
- Replaces primary partition and improves server resiliency

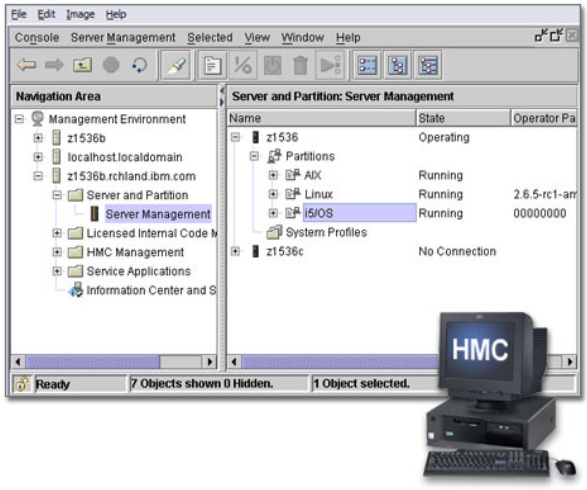

Hardware Management Console (HMC)

The Hardware Management Console (HMC) for IBM eServer is the new unified console for all POWER5 servers. The HMC uses its connection to one or more servers (referred to as managed servers) to perform the following functions:

- Creating and maintaining a multipartitioned environment. This encompasses the command capability to run scripted operations, such as moving resources between partitions or starting and shutting down partitions. The HMC lets you define and support up to 254 logical partitions per server. These partitions can be driven by any combination of OS/400, Linux, or AIX operating systems.
- Displaying a virtual operating system session for each partition, including granular operator authentication and authorization
- Displaying virtual operator panel values for each partition, including remote server and control panel.
- Detecting, reporting, and storing changes in hardware conditions as well as acting as a Service Focal Point and Service Agent by gathering and reporting server error events.
- Powering managed servers on and off.
- Activating Capacity On Demand resources.

The HMC has a GUI interface for many common functions. In addition, On POWER5 eServer i5 servers, the HMC has a 5250 emulator for i5/OS console operations. Virtual serial line interface protocol (virtual SLIP) is also supported, providing a TCP/IP pipe through the service processor and POWER Hypervisor to partitions running under i5/OS. No additional network adapters are required on the server for these functions. On POWER5 eServer p5 and OpenPower servers, many of the functions are performed either through the HMC command-line interface or through a built-in Advanced Systems Management Interface. *(**Note**: For more information on the HMC, see the eServer Hardware Information Center, listed in the "Additional information" section at the end of the IBP.)*

Storage virtualization

- **Storage spaces**
  - 1 MB to 1 TB each
  - Up to 32 per Integrated xSeries Server
  - Up to 64 per Linux partition
  - Can be dynamically added
- **Enables other operating systems to leverage advanced eServer i5 storage architecture**
  - Data automatically spread and protected
  - More disk arms for better performance
  - Automatic balancing of storage across drives

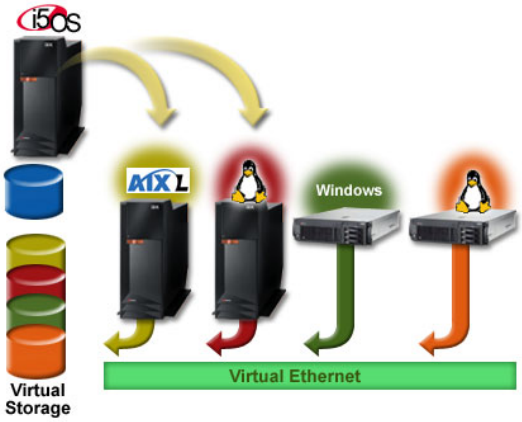

Storage virtualization

Another way that IBM provides better system resource utilization is through its support of storage virtualization. Storage virtualization provides the ability to dedicate I/O adaptors and devices to a virtual server, allowing the allocation and management of I/O devices needed in an On Demand Business environment. With virtual storage, the physical disk drives can be managed by one operating system (for example, i5/OS), while the storage spaces are allocated and linked to another operating system partition (for example, Linux). With storage virtualization, the addition of virtual disk storage spaces to a partition can be performed dynamically, while the partition continues to run. Storage virtualization enables one operating system to leverage the advanced storage architecture of another, and to improve system utilization.

Virtual LANs

- **Provides 1-Gb connections with no LAN adapters/switches**
  - Between partitions and Integrated xSeries Servers
- **Benefits**
  - Extremely secure and reliable server communication over high-performance internal paths
  - Can reduce network traffic and exposure to "sniffing"

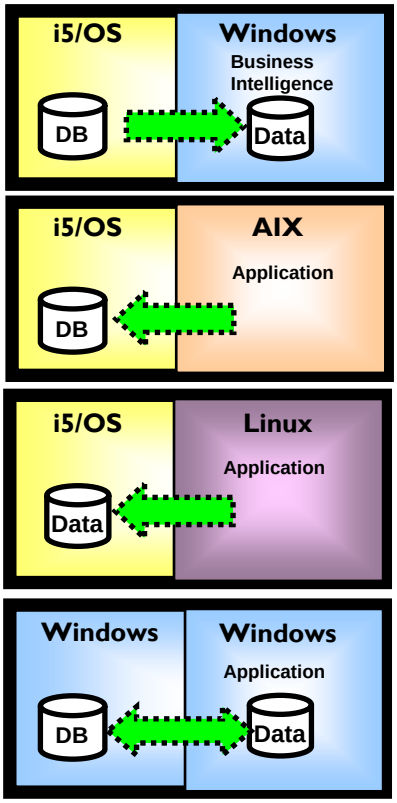

Virtual LANs

IBM virtual local area networks (VLANs) provide high-speed, in-memory network connections between partitions without any dedicated physical I/O slots and network adapters. The POWER Hypervisor recognizes which partitions are located on the same physical server and provides a communications shortcut when these partitions communicate with each other. Partition to partition communications via VLAN (IEEE 802.1Q) are isolated from partitions that are not on the same VLAN, and therefore, providing high degrees of security as well.

Flexible Capacity On Demand options

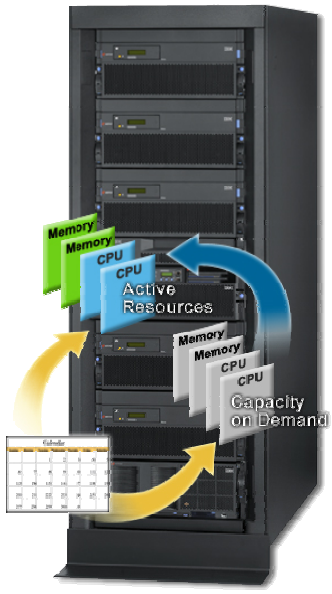

- Extend Capacity Upgrade On Demand (CoD) and On/Off Capacity On Demand options to eServer iSeries and pSeries servers
- Memory CoD
- Trial CoD
- Reserve CoD
  - Automatically enables reserve capacity if processor utilization reaches 100%
- Immediate activation, no system restarts or database reconfiguration required
- *An IBM Virtualization Engine systems technology*

Flexible Capacity On Demand options

A constantly evolving business environment is driving change at an unprecedented pace, which means that a server must be able to adapt at the same rate. IBM eServer systems have exceptional flexibility to adjust dynamically to your business priorities. With Capacity On Demand (CoD), you can turn processors on when you need them and turn them off when you do not. You pay only for what you activate, and not a penny more.

On/Off CoD is perfect for spikes in needed capacity (also known as peak loads). You will be required to pay for each Processor Day you request, but you only request what you need, when you need it. The number of Processor Days is determined by how many processors are temporarily activated over what period of time (“quantity of processors” multiplied by “quantity of days activated”).

A trial period of Capacity On Demand allows you to evaluate, at no charge, the use of inactive processors, memory, or both. Once started, the trial period is available for consecutive 30 power-on days. This means the trial period advances only while the server is powered on. It is not possible to stop and restart the trial period; if the trial period is stopped prior to the expiration of the entitled 30 power-on days, the remaining days are forfeited.

The Reserve CoD option is also perfect for spikes in needed capacity. But unlike On/Off CoD, a prepaid feature, which is purchased upfront, Reserve CoD sets a value on the server that represents the number of Processor Days available for reserved CoD capacity. By paying for the reserve capacity ahead of time, no contracts and no reporting to IBM are required. Reserve CoD allows you to place a quantity of reserve processors into the server's shared processor pool. When the server recognizes that nonreserved processors used in uncapped partitions have been 100% utilized, a Processor Day (good for a 24-hour period) will be subtracted from the available total.

Warranties for all processors (startup and standby) begin when the server is installed. This immediacy of warranty activation allows the enterprise to be better positioned to receive prompt service for the repair and replacement of any failing parts, which helps to ensure that the full complement of standby processors are available for immediate activation.

Maintenance agreement charges are based on the number of processors actually activated; maintenance prices will be adjusted for each permanent processor activation, just as with any other hardware feature addition. (**Note**: Temporary activations do not affect maintenance pricing.)

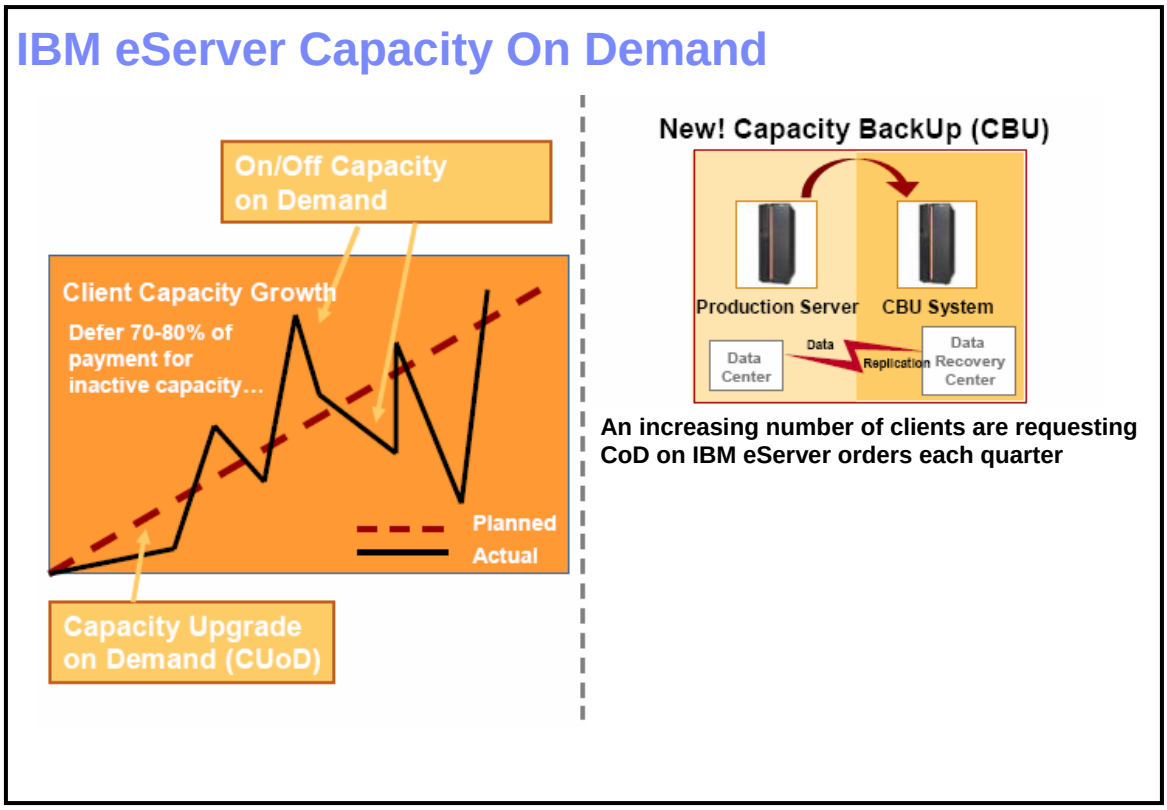

IBM eServer Capacity On Demand

Businesses that use IBM eServer systems are increasingly enthusiastic about the IBM Capacity On Demand offerings.

On/Off Capacity On Demand offers the potential to dramatically improve your ability to react to intermittent or unexpected spikes in user demands for service, without bearing the cost of excess capacity when the workload is less demanding.

Capacity Backup On Demand can enable operations to be sustained on a backup server in the event of a disastrous failure (or planned outage) of a production server.

IBM Director Multiplatform

- Provides centralized management across heterogeneous servers
  - Collect inventory
  - Establish monitors
  - Set alerts
  - Take automatic actions
- Expanded to support more operating systems
  - Server: i5/OS V5R3, Windows, Intel Linux
  - Agents: Servers + POWER Linux, AIX 5L
- Integrated with PM iSeries for collecting and reporting multi-OS CPU utilization and capacity planning
- Complements iSeries Navigator
- IBM Virtualization Engine Systems Service

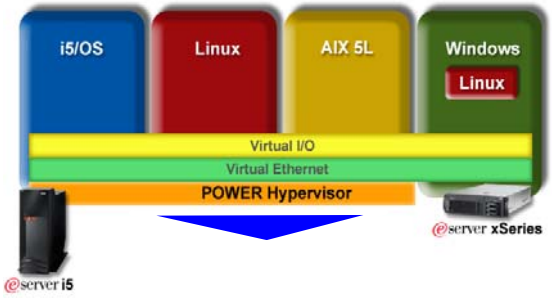

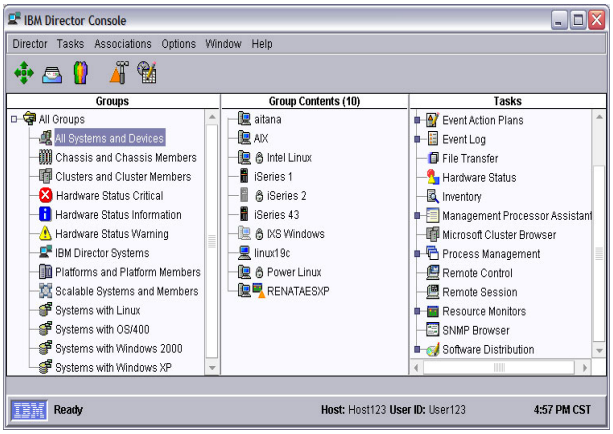

IBM Director Multiplatform

IBM Director Multiplatform is designed to provide a consistent systems management infrastructure so that heterogeneous servers and their resources might be managed together. It exploits proven capabilities from IBM Director that is also delivered on the IBM eServer xSeries® platform for the management of Intel®-based servers.

Coupled with the Virtualization Engine console, which provides a common Web-based console for monitoring and managing the overall health of an On Demand Business environment, IBM Director Multiplatform provides powerful system management capabilities and is based on IBM Director Version 4.2.

IBM Director Multiplatform complements iSeries Navigator and Management Central functions and has similar capabilities, such as hardware and software inventory. It contains system monitors for CPU utilization and threshold triggers, as well as event logs for monitoring across multiple systems or logical partitions.

Now that you have examined the POWER5 common hardware and technology capabilities, let us examine in more detail the function and value of the operating systems supported by the POWER5 servers (AIX, i5/OS, and Linux).

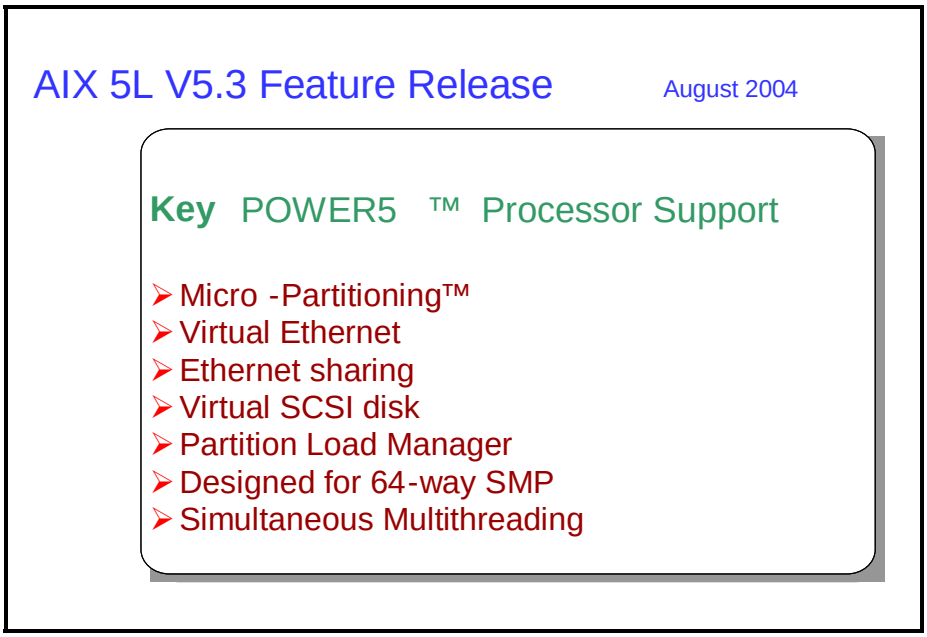

AIX 5L Version 5.3 release content

For more than a decade, businesses large and small have relied on the AIX operating system, which is an industrial-strength UNIX environment, for mission-critical applications. In today's competitive world, businesses need a safe, secure, stable, and flexible operating environment on which to run their organizations. That is why the AIX 5L for POWER operating system is rapidly emerging as the preferred platform for corporate IT managers and independent software vendors. With its proven scalability, reliability, and manageability, AIX 5L is an excellent choice for building a flexible information technology infrastructure. Among all UNIX alternatives, AIX 5L leverages IBM experience in building solutions that run businesses worldwide. And, only one UNIX operating system leads the industry in vision and delivery of advanced support for 64-bit POWER5 platforms, IBM Virtualization Engine systems technologies, and affinity for Linux. That operating system is AIX 5L.

Since August 2004, with the release of Version 5.3, AIX 5L has offered new levels of innovative self-management technologies. It continues to exploit current 64-bit system architecture to support advanced virtualization options, as well as POWER5 processors with SMT capability for improved performance and system utilization.

As with the i5/OS operating system, AIX 5L V5.3 is now enhanced to support the IBM Virtualization Engine technology innovations available on POWER5 systems, including Micro-Partitioning and virtual I/O support. This AIX 5L release underscores firm IBM commitment to deliver business value in a suite of products that employ the latest and most advanced POWER technologies.

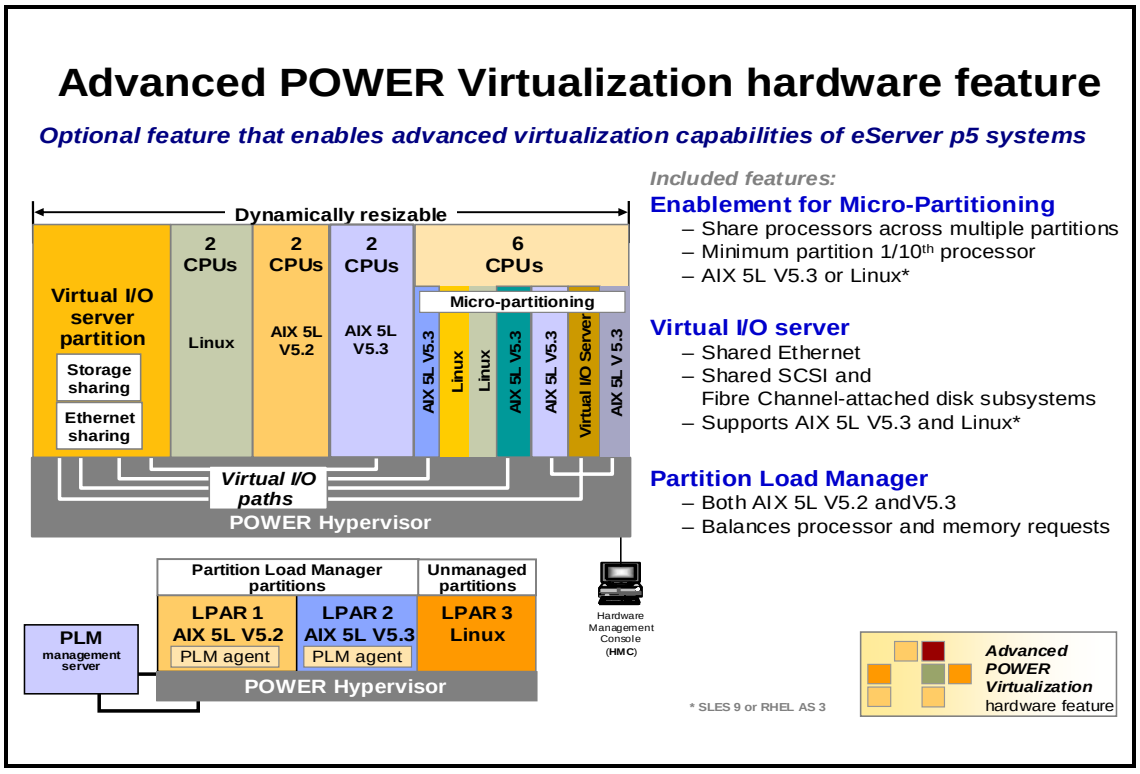

Advanced POWER Virtualization hardware feature

Here is an overview of how AIX 5L applies the *Advanced POWER Virtualization* feature provided in the POWER5 processor hardware. In addition to support of dedicated partitions, AIX V5.3 enables Micro-Partitioning, virtual I/O, virtual Ethernet, virtual storage, and Ethernet sharing. Not only does Micro-Partitioning provide greater granularity in the use of processor resources, it also assures the individual partition's security and software fault isolation. Virtual Ethernet supports in-memory network connections, which provides high-bandwidth network connectivity to all other partitions within the system without having to dedicate physical I/O slots and network adapters. Virtual I/O in AIX 5L V5.3 offers businesses the ability to run partitions with no physical I/O adapters (boot and run from virtual SCSI, network connections via virtual Ethernet, and Ethernet adapter sharing).

To meet the ever-changing demands of today's business operations, nearly all of these features are dynamically reconfigurable. IBM understands that monitoring the rapid turns of the business world, and simultaneously managing every aspect of the system's performance is a complicated and resource-consuming task. AIX V5.3 not only provides system administrators with the flexibility of dynamic resource allocation, but it also equips them with a feature called Partition Load Manager, which automatically balances system loads based on a set of predefined parameters.

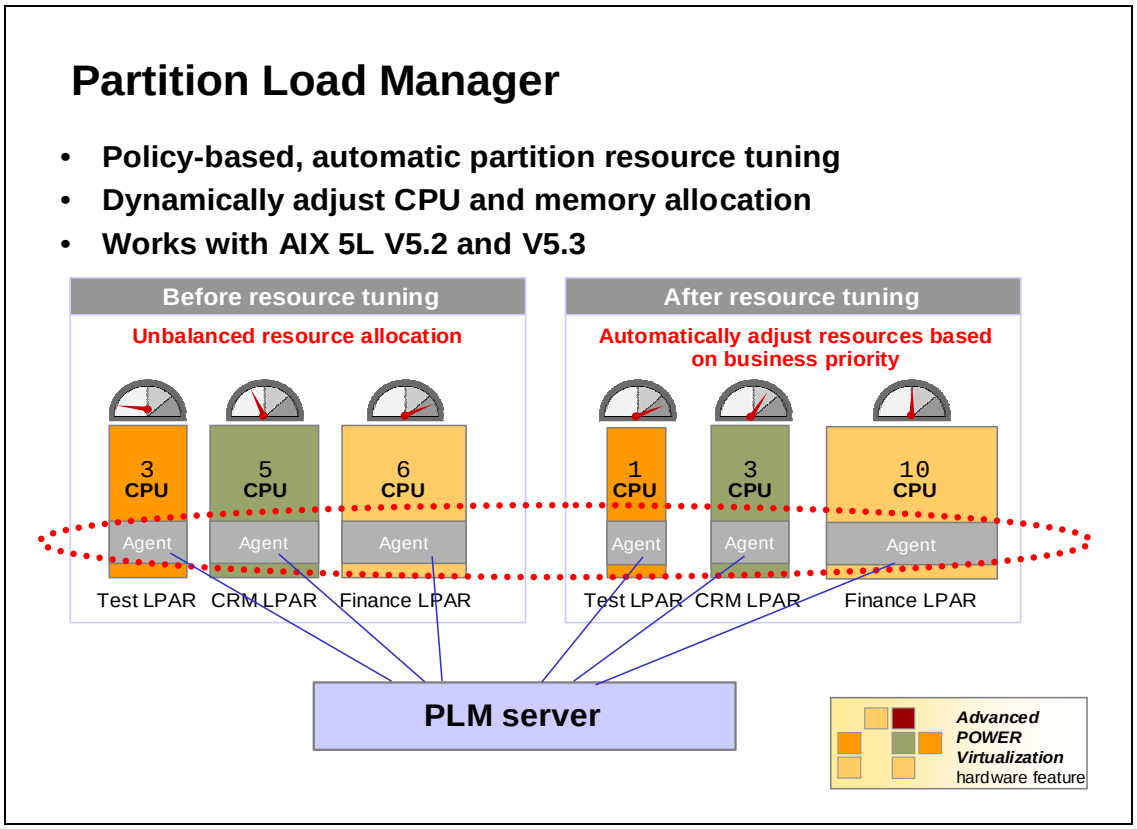

Partition Load Manager

The Partition Load Manager (PLM) software is part of the Advanced POWER Virtualization feature. It helps businesses maximize the utilization of processor and memory resources of DLPAR-capable logical partitions running AIX 5L on eServer p5 servers.

The PLM is a resource manager that automatically orchestrates the movement of processor and memory resources between partitions based on defined policies and utilization of these resources. PLM manages memory as well as dedicated and shared processor partitions using Micro-Partitioning technology to readjust the resources. PLM adds greater flexibility to the Micro-Partitioning function offered by the POWER Hypervisor.

Business policies might specify the partitions that run higher priority tasks. When there are increased needs for computing power in these partitions, PLM can take away resources from lower priority workloads and automatically redeploy them to high business value tasks. In many cases, this can significantly improve the effectiveness of business operations. Because PLM automatically monitors the load in each managed partition and tunes the processing and memory configuration without operator intervention, it also helps corporations reduce the overall cost of system administration.

IBM eServer i5 and iSeries with i5/OS V5R3

Builds on and extends the capabilities of OS/400 V5R2

The world's most complete, secure, integrated system designed to run thousands of business applications and simplify the IT environment so you can save money and reinvest in growing the business.

IBM eServer i5 and iSeries with i5/OS V5R3

The IBM eServer i5 and iSeries family of servers and its integrated i5/OS operating system has grown into one of the most popular midmarket server franchises available today. More than 400,000 iSeries and AS/400 systems have been sold in over 100 countries around the world.

i5/OS V5R3 is the current operating system release for eServer i5 and iSeries systems. The history of the iSeries systems and i5/OS goes back to 1988 with the announcement of the IBM AS/400 server line and its accompanying operating system, OS/400. All releases of i5/OS released prior to V5R3 were called OS/400, and thus, the terms "OS/400" and "i5/OS" are used interchangeably when referring to the integrated iSeries operating system.

What makes the i5/OS and OS/400 operating systems unique is that they provide the most complete and secure, integrated operating system that is designed to run thousands of business applications. Enterprises can incorporate new technology and function while minimizing disruption, and enjoy a simplified IT environment that saves time and money.

This value proposition unfolds to IBM clients and solution providers through many examples of nondisruptive change. Consider the fact that IBM has improved the underlying iSeries hardware components many times, including the processor technology, bus architecture, I/O capabilities, and supported devices, yet the business applications running under OS/400 and i5/OS have seamlessly transitioned from one release to the next. In nearly all cases, a smooth transition has occurred without even requiring recompilation of the applications. Yet with each hardware and operating

system innovation, these applications are able to exploit the incremental function and new capabilities provided by IBM.

Perhaps one of the best examples of the phenomenal ability for iSeries systems and i5/OS to adapt to new technologies was demonstrated when the iSeries hardware evolved from 48-bit CISC into 64-bit RISC processor technology. On most other platforms, this significant change of processor technology would have required a monumental programming, recompilation, and testing effort. However, the iSeries platform handled this change to server hardware as it had in the past, with minimal, if any, impact to users. Because of this unique architecture, thousands of iSeries business applications were automatically translated into 64-bit applications without any programmer intervention.

Insulating businesses and application developers from continuous improvement in the underlying iSeries hardware is a hallmark of i5/OS. Let us examine i5/OS in more detail to explain how this is possible.

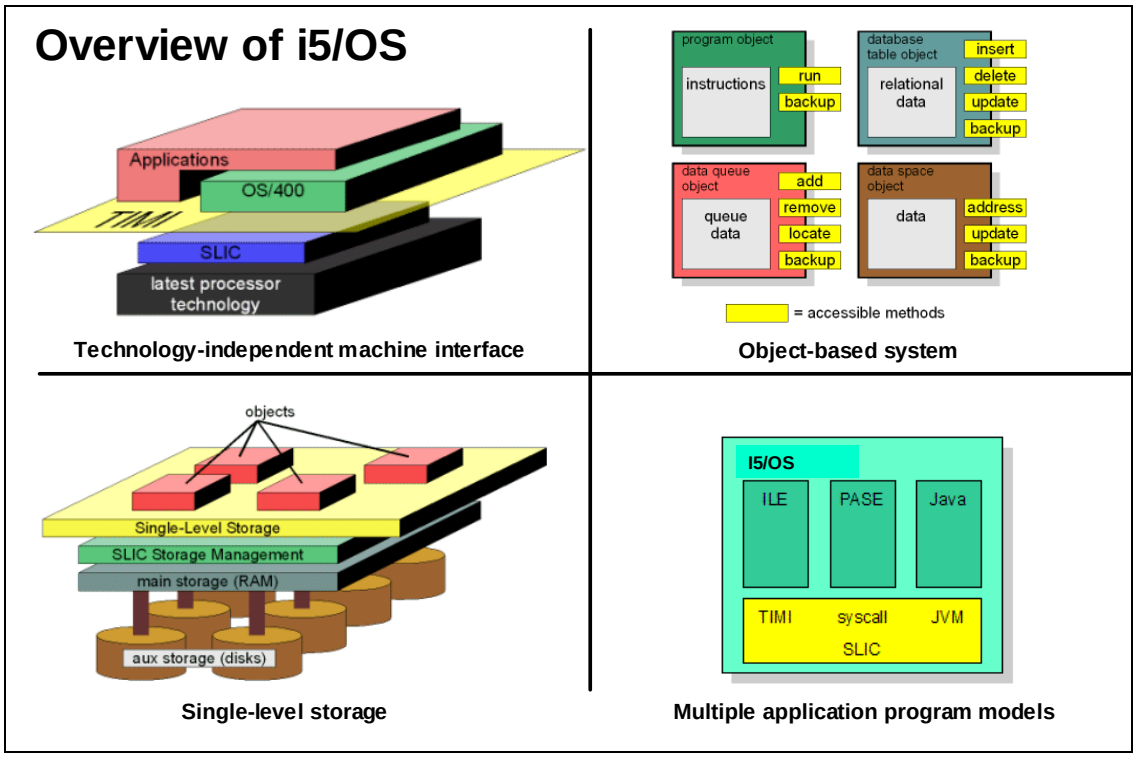

Overview of i5/OS architecture

Several features of the i5/OS architecture distinguish the system from other machines in the computing industry. These features include the following:

- Two-part primary operating system
- Technology-independent machine interface (TIMI)
- Object-based system
- Single-level storage
- High degree of integration
- Multiple application program models
- Open standards

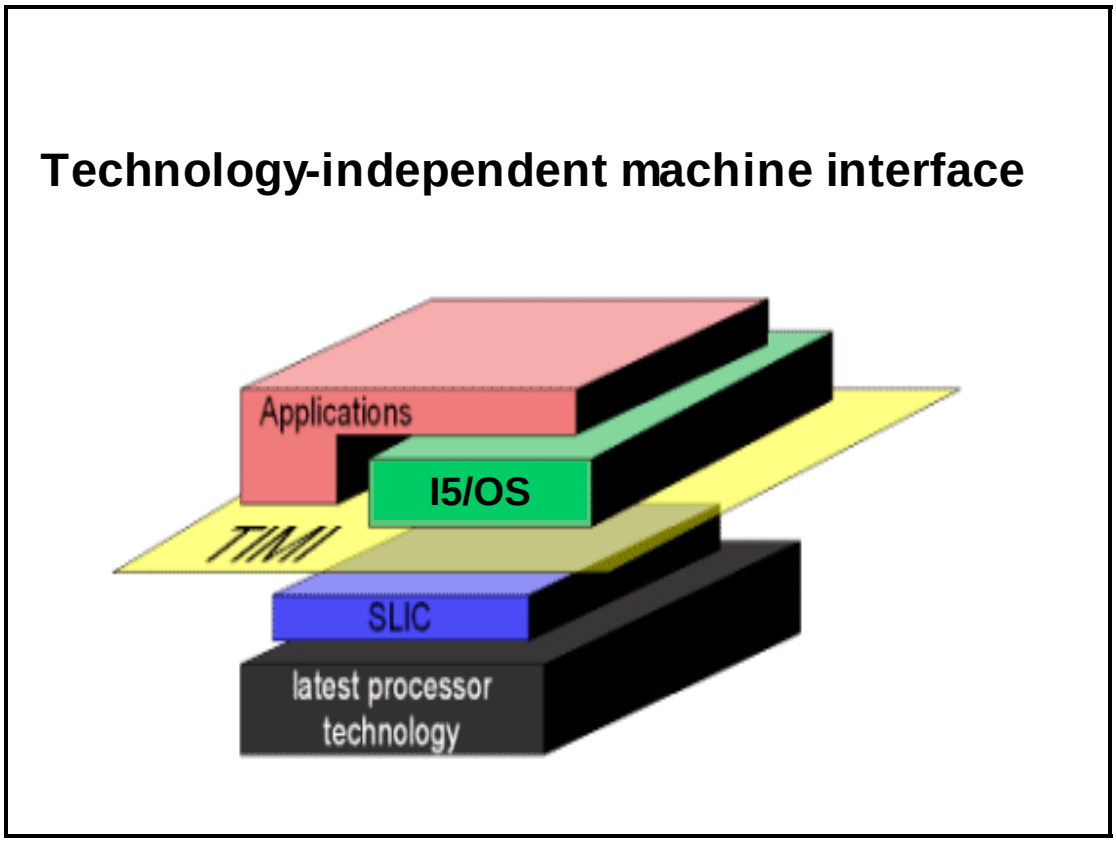

Technology-independent machine interface

There are two components to the operating system software on an iSeries server. This important distinction is unique in the industry in its completeness of implementation. The two components are System Licensed Internal Code (SLIC) and i5/OS.

SLIC provides: the technology-independent machine interface (TIMI), process control, resource management, integrated SQL database, security enforcement, network communications, file systems, storage management, JVM, and other primitives. SLIC is a hardened, high-performance layer of software at the lowest level, much like a UNIX kernel, only far more functional.

i5/OS provides higher-level functions based on these services to users and applications. i5/OS also delivers a vast range of high-level language (such as C/C++, COBOL, RPG, and FORTRAN) run-time functions. i5/OS interacts with the client-server graphical user interface, iSeries Navigator.

At a macro level, a logical partition (LPAR) running the traditional iSeries operating system, can be referred to as running i5/OS. The name i5/OS can refer to either the combination of both parts of the operating system or just the "top" portion.

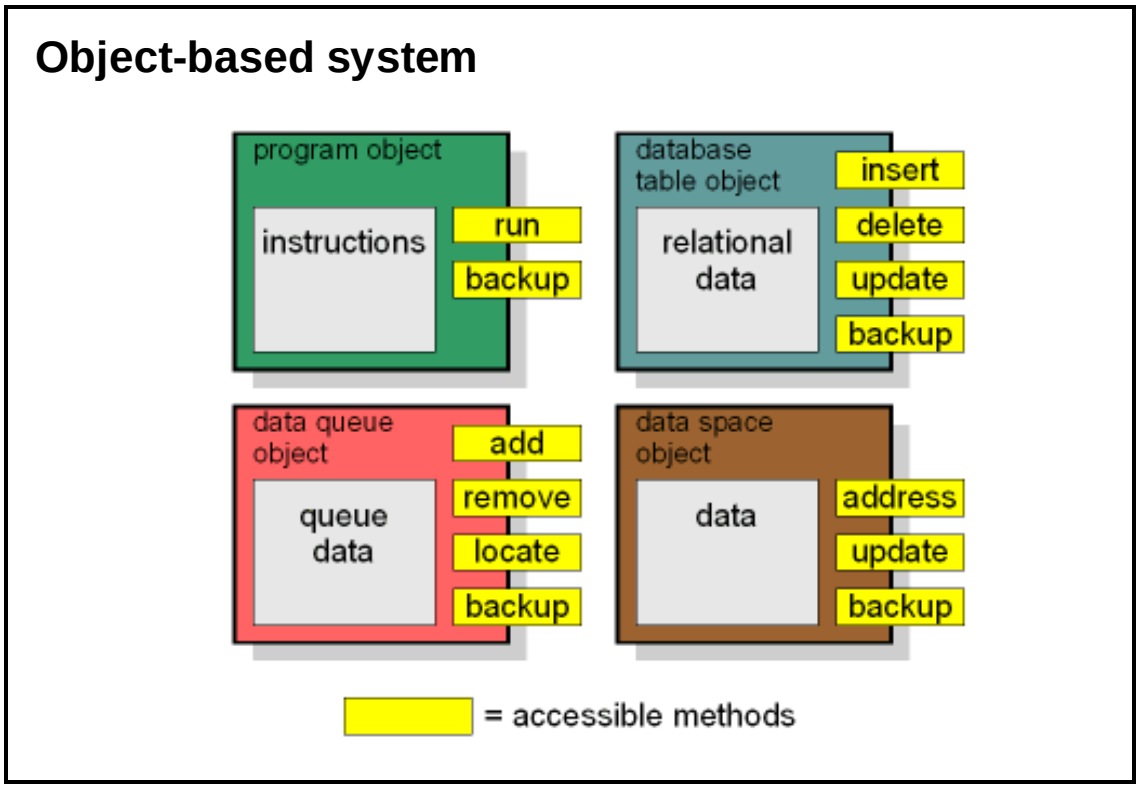

Object-based system

i5/OS maintains all information as objects. There are hundreds of object types, including the four examples above (program, database table, data queue, and data space). This is different from the simple byte-string, file-based manipulation used by many systems. Object-based design enables a powerful, yet manageable level of system integrity, reliability, and authorization constraints.

All programs and operating system information, such as user profiles, database files, programs, and printer queues, have their associated object types stored with the information. In the i5/OS architecture, the object type determines how the information contained in the object can be used (that is, by which methods the information can be accessed).

For example, it is impossible to corrupt a program object by modifying its code sequence data as if it were a file. Because the system knows the object is a program, it will only allow valid program operations (for example, run or backup). Thus, with no write method, iSeries program objects are, by design, highly virus-resistant. Other kinds of objects include directories and simple stream data files, such as video and audio files. These stream-file objects provide familiar open, read, and write operations.

Single-level storage

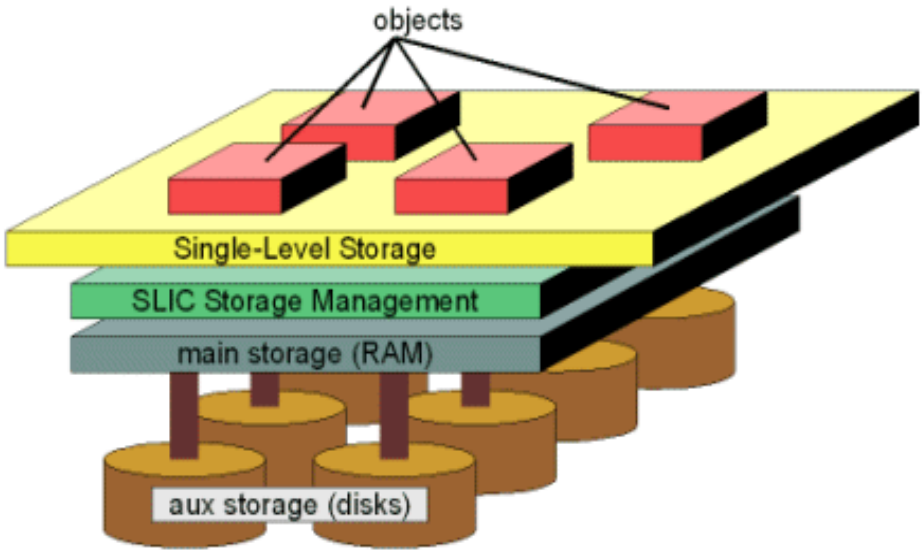

Single-level storage

i5/OS applications, and the objects with which they interact, all reside in a very large virtualized, single-level storage. That is, the entire system, including the objects most other systems distinguish as "on disk" or "in memory," are all in the single-level storage. Objects are designated as either permanent or temporary. Permanent objects exist across system initial program loads (IPLs) (also known as reboots). Temporary objects do not require such persistence. Essentially, the physical RAM on the server is a cache for this very large, single-level storage space. Storage management, a component of SLIC, ensures that objects needing to persist when the system is off are maintained in persistent storage. This is either magnetic hard disk or flash memory.

The benefit of providing a single, very large address space, in which all objects on the system reside, is that applications need not tailor their memory usage to a specific machine configuration. In fact, because of the single-level storage, i5/OS programmers do not have to tailor such things as the disk cache size or the paging space. This greatly facilitates the on demand allocation of memory among logical partitions.

Multiple application program models

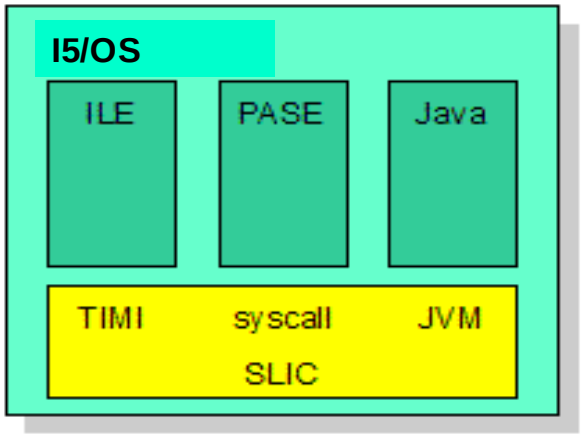

Multiple application program models

The traditional programming model for i5/OS applications uses the IBM Integrated Language Environment® (ILE). i5/OS provides many programming interfaces to interact with system resources, most of which are ILE interfaces. Applications in this model benefit from the TIMI for complete protection and benefit from future hardware changes. Support is provided for C/C++, COBOL, RPG, and CL.

Applications from the AIX operating system can run inside an i5/OS job, completely and seamlessly integrated into i5/OS. This option is the i5/OS Portable Application Solution Environment (PASE). Because the microprocessors are the same between traditional AIX and i5/OS hardware, many components from AIX can run directly "on the silicon" in an iSeries system. PASE provides this very easy way to run, often unmodified, AIX applications on an iSeries system. PASE applications, just like AIX applications, interact with operating system functions through a syscall interface. However, because the applications are compiled to a specific processor instruction set, the investment protection in application code offered by the TIMI is not afforded PASE applications. Support is provided for C/C++, FORTRAN, and COBOL.

Support for Java™ applications is also built into i5/OS. Java applications utilize a JVM™ and Java compiler that are built into the i5/OS kernel (SLIC). The object-oriented nature of Java is an ideal fit for an object-based operating system such as i5/OS. If the Java application uses a Web application server, then on i5/OS, it will best utilize WebSphere Application Server - Express, included with every iSeries system.

All three of these application program models normally coexist on i5/OS and run on the same set of system resources (network communications, user security, and file systems). It is common for code in one of these program models to call into code in another. The combination of these models affords great flexibility in bringing new applications to i5/OS.

eServer i5 and iSeries with i5/OS architecture summary

Fostering growth and innovation in solutions
There is real value in integration

- Higher reliability
- Simplicity
- Optimization
- Automation
- Security
- Ease of operations
- Lower overall TCO

The architecture securely paves the way for technology leadership without disrupting your business and has the versatility to respond to opportunities and grow your business.

ibm.com/eserver/iseries

IBM eServer i5 and iSeries with i5/OS architecture summary

The eServer i5 and iSeries servers, coupled with the i5/OS architecture are different than all other architecture in the computing industry. It is flexible and entirely focused on business computing. It is largely self-managed, so the enterprise can focus on running the business, not the computer. It is based on the belief that a radically better system can be designed and built. The unique architectural foundation of the eServer i5 and iSeries platform with i5/OS will continue to allow leadership in technological advancements well into the future.

Linux on POWER

- **Linux brings:**
  - Open standards
  - Flexibility
- **POWER brings:**
  - Availability and reliability
- **Many distributions (built on Linux 2.6 kernel) for Power Architecture**
  - SUSE Linux Enterprise Server (SLES) V9
  - Red Hat Enterprise Linux Advanced Server (RHEL AS) V4

Linux on POWER

Linux brings open standards, along with maximum availability and flexibility, to your business solutions. Add the proven performance and reliability of the POWER processor-based IBM eServer family of servers, including the IBM eServer OpenPower™ servers, built specifically for Linux, and applications will reap the combined advantages of Linux on POWER.

There are many Linux distributions that run on the POWER and IBM PowerPC® architecture. IBM currently provides support for the following distributions on its POWER processor-based servers:

- SUSE™ LINUX Enterprise Server (SLES): Version 9 is the latest release of SLES and is referred to as SLES9 throughout the rest of this document.
- Red Hat® Enterprise Linux Advanced Server (RHEL AS): Version 4 is the latest release of RHEL AS and is referred to as RHEL4 throughout the rest of this document.

The latest versions of both distributions are built on the Linux 2.6 kernel, which features many functional improvements over its earlier 2.4 kernel version and also enables many of the new POWER5 virtualization features described in previous charts.

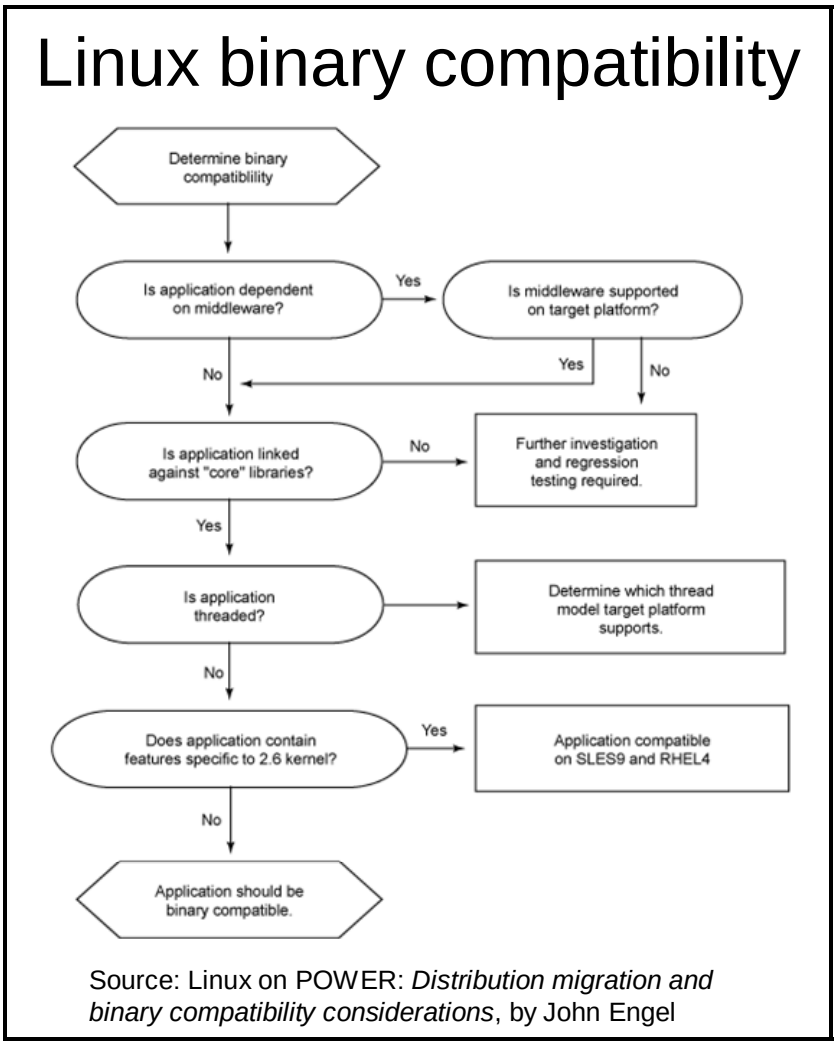

Source: Linux on POWER: *Distribution migration and binary compatibility considerations*, by John Engel

Linux binary compatibility

Binary compatibility refers to the ability to run a binary piece of Linux code on multiple Linux distributions for a given processor family. These environments can be different versions of the same distribution or they can be distinct distributions. For example, a binary that was compiled and executed on a POWER4 processor-based system running SLES9 can also execute on a POWER5 system running SLES9. Similarly, a binary that was compiled and executed on a POWER4 system running RHEL3 can also execute on a POWER5 system running RHEL3.

From the standpoint of the hardware instruction sets, code compiled on POWER4 hardware has a very high probability of running unchanged on POWER5 and PowerPC 970(FX) hardware. However, as you can see from the flow chart shown here, it is really the supported levels of the Linux GCC (GNU Compiler Collection) compilers, libc (the C library), the middleware, Java, and the Linux distributions themselves that ultimately determine binary compatibility for applications between hardware versions. For example, if an installed distribution was exactly the same on POWER4 hardware and POWER5 hardware, an application compiled on either hardware platform is compatible.

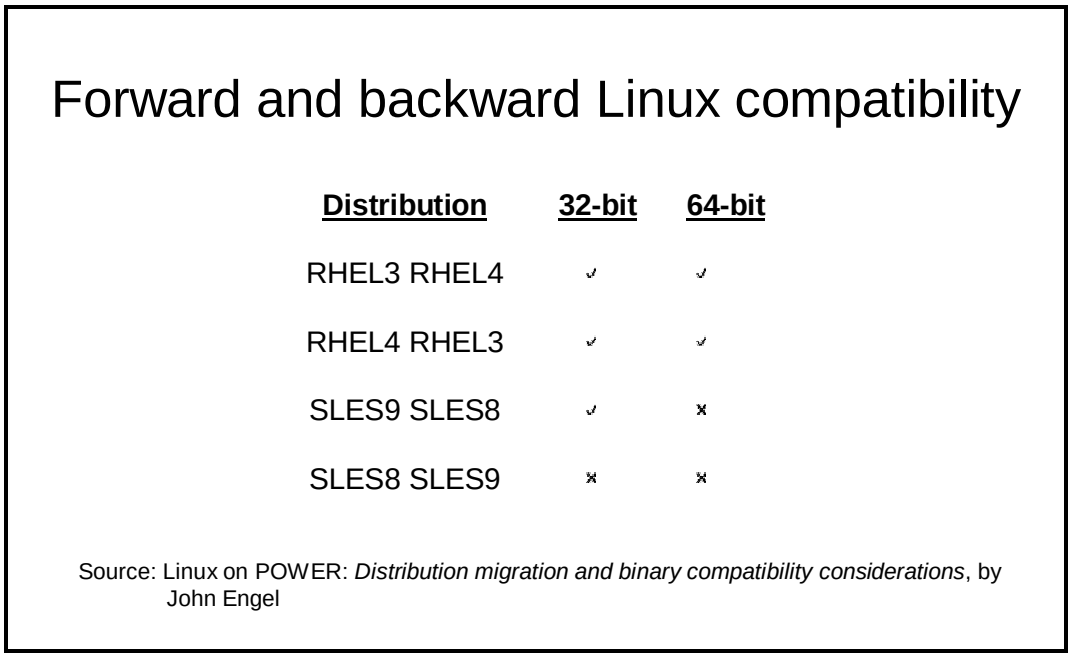

Forward and backward Linux compatibility

Red Hat supports 32-bit and 64-bit forward and backward compatibility between RHEL3 and RHEL4, while SLES8 only supports 32-bit backward compatibility when migrating from SLES8 to SLES9. This means that applications developed on SLES8 for POWER processor-based systems in 32-bit mode should require only moderate regression testing to confirm compatibility when deploying on SLES9. For all 64-bit applications developed on SLES8, recompilation and regression testing are required when moving to SLES9. The table shown here summarizes the forward and backward compatibility between RHEL3 and RHEL4 and between SLES8 and SLES9.

Both 32-bit and 64-bit addressing is fully supported by the same 64-bit Linux kernel on eServer OpenPower, eServer p5, eServer i5, and IBM eServer BladeCenter® JS20 servers. Moreover, no performance penalty is to be expected when executing 32-bit applications on 64-bit systems. POWER5 and PowerPC 970(FX) chips support the two addressing types directly. Both SLES9 and RHEL4 are based on the Linux 2.6 kernel.

(**Note:** For more information, read the article *Linux on POWER: Distribution migration and binary compatibility considerations* found at the link provided in the “Additional information” section of this course.)

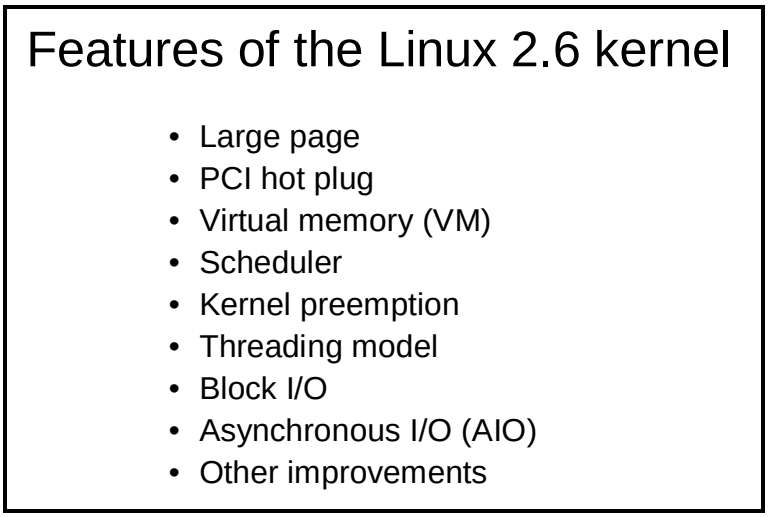

Features of the Linux 2.6 kernel

The 2.6 kernel is the latest evolution of the Linux kernel, offering many new features:

Large page: The 2.6 kernel supports two virtual page sizes: the traditional 4-KB page size and the 16-MB page size (for memory-intensive applications).

PCI hot plug: This function allows you to insert a new PCI hot plug adapter into an available PCI slot while the operating system is running.

Virtual memory (VM): A new and efficient algorithm, called reverse mapping, for virtual-to-physical page mapping that enables more intelligent memory swapping.

Scheduler: A new scheduler improves load balancing and the performance of interactive applications.

Kernel preemption: In the 2.6 kernel, a kernel task can be pre-empted. This not only improves performance for user-interactive applications, but also benefits embedded devices and real-time systems.

Threading model: The improved threading model, Native POSIX Thread Library (NPTL), associates one kernel thread with one user thread to support multithreaded enterprise-level applications, as well as Web server and Java applications.

Block I/O: The newly rewritten block I/O improves performance and error handling.

Asynchronous I/O (AIO): AIO is now supported at the kernel level (instead of the user level) and runs in the background without blocking user tasks, thus improving performance.

Other Improvements: There are other improvements to the 2.6 kernel. The number of users and groups has been increased from 65,000 to over four billion. The number of PIDs has been increased from 32,000 to one billion. Support for 16-TB file systems. The limit on the number of open files no longer needs to be set in advance.

Compilers for Linux on POWER

- IBM XL C/C++
  - -qarch and –qtune for optimization
  - -qaltivec for the VMX feature
  - Produces compilations that utilize more POWER5 advanced hardware features
- IBM XL Fortran
- GNU Compiler Collection (GCC)
  - -mcpu-power5 and –mtune=power5 for POWER5-specific instructions
  - VMX vector extensions
  - Produces more portable compilations

Compilers

The compilers shown on this screen are available for Linux on POWER:

The high performance compiler, IBM XL C/C++ Version 7.0, is available for RHEL4 and SLES9 and adds performance improvements for POWER5 processor-based systems. The **-qarch** and **-qtune** options optimize the environment for their respective architectures. For example, to optimize for POWER5, use these options: **-qarch=pwr5** and **-qtune=pwr5**. The **-qaltivec** option in the IBM XL C/C++ compiler enables the VMX feature on the PowerPC 970 and 970FX processors in IBM eServer BladeCenter JS20 servers.

The GNU Compiler Collection contains compilers for different programming languages. Many improvements were made from Version 3.2 to 3.3, including POWER processor-specific optimizations to its C compiler, GCC. The **-mcpu=power5** and **-mtune=power5** flags are now supported. They result in instructions that are specific to the POWER5 architecture. There are also VMX vector extensions for the IBM PowerPC 970 and 970FX processor that can increase the performance of vectorized code.

Because the IBM XL C/C++ compiler uses more POWER5 advanced hardware features, it provides better performance for the programs compiled with it, whereas programs compiled using GCC are more portable.

(**Note**: More information about using the XL C/C++ compiler for Linux on POWER can be found in the article "How to use IBM XL C/C++ Advanced Edition V7.0 for Linux on POWER: A guide for GCC users." A link to this site is found in the “Additional information” section of this course.)

Overview of Linux on POWER hardware

- IBM eServer OpenPower systems
- IBM eServer p5 and eServer pSeries systems
- IBM eServer i5 and eServer iSeries systems
- IBM eServer BladeCenter JS20

Overview of Linux on POWER hardware

The IBM eServer OpenPower, IBM eServer p5, and IBM eServer i5 servers are based on the same POWER5 architecture.

Linux is the only operating system supported on eServer OpenPower systems. The two currently available models (OpenPower 710 and OpenPower 720) target businesses that are making a long-term commitment to Linux and those desiring leadership price and performance with enterprise-class reliability and virtualization characteristics.

IBM eServer p5 and pSeries systems are intended for enterprise class UNIX users and for those seeking the most scalable and reliable Linux and AIX servers. In addition, support for IBM i5/OS is available in some models of the eServer p5 product line.

IBM eServer i5 and iSeries systems offer an on demand computing environment for IBM i5/OS, AIX 5L, Microsoft Windows, and Linux solutions. These highly integrated servers reduce complexity and enhance productivity through server consolidation.

eServer BladeCenter JS20 offers a high density footprint of servers. Up to 84 2-way blades can be installed in a 42U rack, which reduces space and cooling requirements. Its modular design supports secure scalability in a rapid and cost effective manner. Virtualization, which was discussed earlier in this course, is not available on eServer BladeCenter JS20.

Linux adds to the POWER5 flexibility message

- Linux on POWER provides a proven, open, and powerful computing platform.
- IBM offers powerful support for Linux application developers.
- Migration from Linux on Intel to Linux on POWER brings flexibility, scalability, and out-of-the-box innovation.

Linux adds to the POWER5 flexibility message

In the world of computing, demand is growing for flexibility. In the past, the pursuit of flexibility exacted many trade-offs, most notably in price and performance. But with the advent of Linux and 64-bit processors, support for which is delivered on the POWER5 architecture, the trade-offs are quickly fading away. Now you can run almost any application, anytime, anywhere, without taking a price or performance hit.

IBM is committed to leveraging Linux on POWER to provide the proven, open, and powerful computing platform for you to drive innovation, leap frog your competition, and create value for your company, customers, and shareholders. These high-performance offerings are surprisingly affordable and provide leading-edge reliability, availability and serviceability. This is why Linux on POWER is building such momentum in the industry.

Solution providers are porting their Linux solutions to POWER hardware and businesses are adopting the POWER platform for enterprise-critical applications. And with IBM's public dedication to Linux and the Power Architecture, developers need not go it alone with their Linux initiatives. The IBM Linux Technology Center (LTC) is the premier organization for enterprise Linux development in the world today, and is driven by a major emphasis on POWER solution enablement. This support, combined with traditional support from the Open Source community provides an unprecedented network for developer assistance.

The field is primed for migration from Linux on Intel to Linux on POWER. The motivation to do this is driven by the need of solutions to offer maximum flexibility and scalability, coupled with out-of-the-box innovation.

If you can dream it, Linux on POWER can support it.

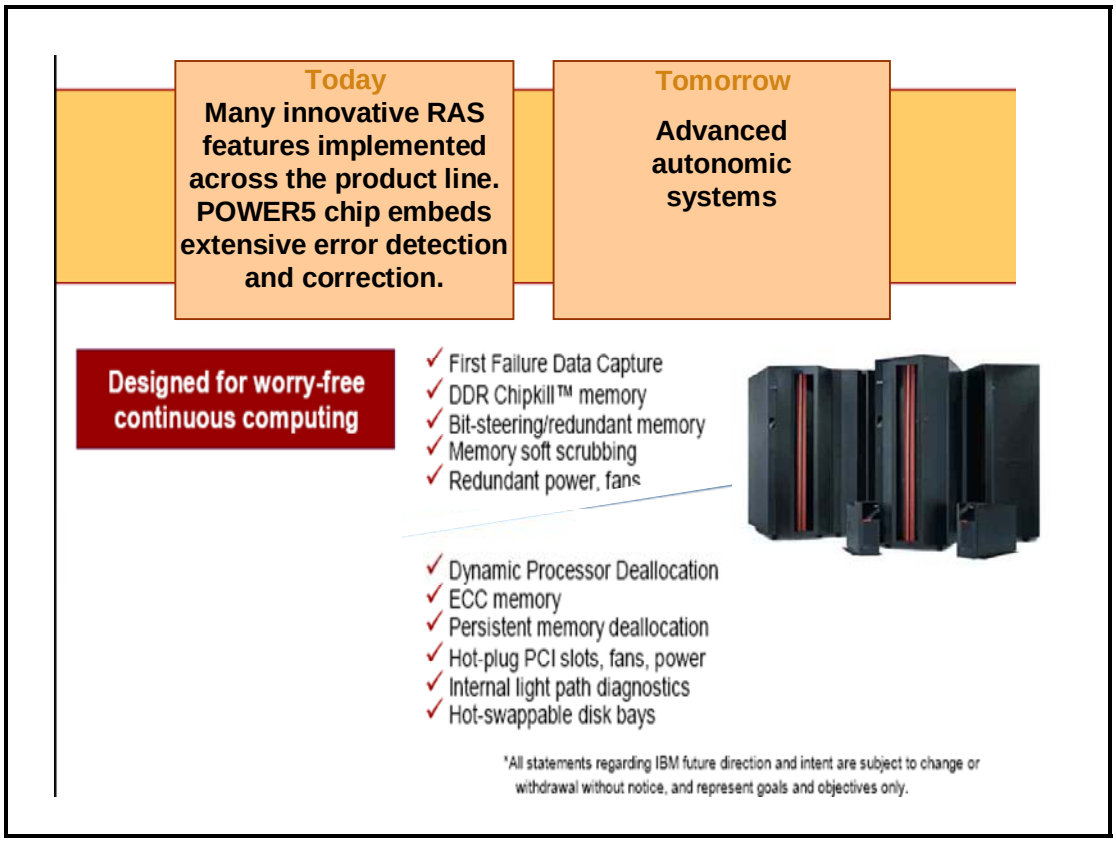

Reliability

Many of the reliability and availability characteristics found in IBM eServer systems are inspired by a long history of mainframe success. Innovative serviceability features, such as light path diagnostics, simplify the process of servicing installed systems.

This list of check-marked features is included across the entire product line, making IBM eServer systems self-managing. The industry refers to this as autonomic management. These capabilities ensure high availability. For example, IBM Peripheral Component Interconnect (PCI) “retry” logic can help prevent server outages caused by parity error detections. The lack of this capability on Sun™ servers can result in unplanned system outages.

The roadmap for tomorrow’s IBM eServer platforms includes even more advanced autonomic (self-managing) systems. The next generation of POWER processors will continue to build on the error detection and correction technology embedded in its current design.

The benefit is a server that is literally designed for worry-free processing, which means you can focus on your business and not wonder when your technology is going to fail.

Summary: Common IBM eServer Platform

IBM eServer pSeries	**Make no compromises, accept no limits in an on demand world** Feature innovative and affordable POWER servers to handle the multiple workloads of the UNIX and Linux marketplace
IBM eServer i5	**Simplicity in an on demand world** Feature fully integrated solutions including IBM middleware with maximum flexibility in operating system choice

IBM eServer POWER5 Common Platform

POWER5 servers are the industry's most powerful and flexible servers, running multiple operating environments simultaneously and dynamically adjusting to changing requirements. Exactly what you need to simplify IT infrastructure, drive down costs and drive up productivity in today's on demand world.

Summary: common IBM eServer platform

In summary, IBM is delivering on its promise to provide state-of-the-art technology and servers. The Power Architecture, together with the IBM industry-leading UNIX and integrated servers, provides a winning combination for today's demanding enterprises. Delivering greater performance, quality, and flexibility in a package that is easier to manage and support.

Additional information

You might want to visit these sites to find out more about the information provided in this course:

- eServer i5 three-in-one POWER5-based benchmark
  ibm.com/servers/eserver/iseries/hardware/threeinone

- Simultaneous multithreading (SMT) on iSeries POWER5 processors
  ibm.com/servers/eserver/iseries/perfmgmt/pdf/SMT.pdf

- Introducing IBM eServer i5 570 and 520
  ibm.com/eserver/iseries

- IBM Power Architecture
  ibm.com/technology/power

- Linux on POWER: Introducing IBM eServer i5
  ibm.com/eserver/linux/power/features/iseries.html

- iSeries Innovations: POWER5
  ibm.com/eserver/iseries/about/innovations.html

- Blue Gene/L supercomputer
  ibm.com/technology/ourwork/casestudies/bluegene.shtml

- IBM Virtualization Engine
  ibm.com/servers/eserver/about/virtualization/index.html

- IBM eServer Software Information Center
  http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/esmcinfo/eicacoverview.htm

- IBM eServer Hardware Information Center
  http://publib.boulder.ibm.com/infocenter/iseries/v1r2s/en_US/index.htm?info/icmain.htm

- Linux article: *Linux on POWER: Distribution migration and binary compatibility considerations*
  ibm.com/developerworks/eserver/library/es-bincomp/

- pSeries Web site
  ibm.com/eserver/pseries

- eServer i5 Web site
  ibm.com/eserver/i5

- iSeries Web site
  ibm.com/eserver/iseries

Trademarks and special notices

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

The following are trademarks or registered trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM, the IBM logo, BladeCenter, eServer, DB2, DB2 UDB, iSeries, xSeries, pSeries, AS/400, OS/400, i5/OS, the i5/OS logo, POWER, Power Architecture, POWER Hypervisor, POWER4, POWER5, PowerPC, OpenPower, RS/6000, TotalStorage, AIX, AIX 5L, Tivoli, WebSphere, On Demand Business, Micro-Partitioning, Redbooks, Integrated Language Environment, Power Family, Chipkill, and ibm.com.

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

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

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

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

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

All other products may be trademarks or registered trademarks of their respective companies.

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

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.