

Commonality Across IBM eServer iSeries and pSeries Servers Running Linux



Kay A. Tate

IBM eServer Enablement

Special Notices

This publication/presentation was produced in the United States. IBM may not offer the products, programs, services or features discussed herein in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products, programs, services, and features available in your area. Any reference to an IBM product, program, service, or feature is not intended to state or imply that only IBM's product, program, service, or feature may be used. Any functionally equivalent product, program, service, or feature that does not infringe on IBM's intellectual property rights may be used instead.

Information in this presentation concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.


Information is provided "AS IS" without warranty of any kind. The information contained in this document represents the current views of IBM on the issues discussed as of the date of publication. IBM cannot guarantee the accuracy of any information presented after the date of publication.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

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.

Some information in this presentation addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

AS/400	IBM	pSeries	POWER4
AS/400e	IBM (logo)	RS/6000	
eServer	iSeries	AIX	
 eServer	OS/400	VisualAge	

Lotus and SmartSuite are trademarks of Lotus Development Corporation and/or IBM Corporation in the United States, other countries, or both.

MMX, Pentium, and ProShare are trademarks or registered trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft and Windows NT are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

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

SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

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

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


Abstract

Moving a solution to Linux on iSeries and pSeries servers that employ the POWER architecture is attractive to many independent software vendors (ISVs). The rapid rise of Linux system use in the marketplace coupled with the “industrial strength” and performance of IBM eServer iSeries and pSeries servers provides a business opportunity for many types of applications. This paper is composed of two parts. The first is a discussion of commonalities in delivering a solution using Linux on iSeries and pSeries servers. The fact that an ISV can do one “port” and deliver product on both of these servers should enhance the return on investment calculation for new platform support. The consolidation of Linux with other operating systems on these platforms provides the ISV with an opportunity to participate in other ISVs’ solutions as well as the opportunity to readily deliver more complete, higher value, solutions on these IBM servers.

The second part of the paper covers a discussion of how an ISV can go about doing this very effectively. A good strategy for delivering new solutions on POWER servers is to select one for primary development and then to do validation testing on the other one. IBM provides programs to assist ISVs with their commonality testing.

Part 1 - Applying POWER Technology to A Solution

Exploiting IBM’s POWER chip technology to deliver Linux applications can give both ISVs and their customers unique and powerful advantages in the marketplace. The advantages start with the chip set itself. IBM POWER technology is built upon a RISC architecture that IBM has delivered in its systems since the early 1990s. From a solid original design, this chip family has been evolved and extended to deliver increased density, parallel operations and overall integration. These extensions include larger numbers of processing units, more and larger and faster caches, more logic in the same package, and higher order logic processing like branch prediction. Integration among system components as well as density has been increased over the chip generations resulting in servers that have been able to deliver rapidly increasing performance and reliability characteristics. IBM has publicly discussed its intent to continue development of this chip architecture family.



Power4 Multichip Module

The performance success of these servers is shown by IBM's notable leadership in public benchmark results. The reliability of the chip set coupled with the system software, can be seen in a July 2003 IDC report on iSeries consolidation. This report shows an order of magnitude less downtime for the defined iSeries configuration running the same workload as a set of other servers. Since POWER technology is a family of processors with the same base architecture, IBM can deliver a high level of compatibility from one generation of the chip set to the next, typically providing upward compatibility for ISVs as they move forward with their solutions. It also provides enhanced computing compatibility to their customers as they move forward in their business process automation, e-business conversion and growth.

POWER technology is the foundation of both IBM eServer iSeries and pSeries systems. With the other system architecture parts IBM has put around the POWER chip sets, these two platforms have demonstrated strength and synergy with applications in both business and high performance computing. While the customer bases for these types of computing overlap, they have some fairly different requirements. POWER technology addresses these set of requirements very effectively. Though the two series have different perceived strengths, the use of a common chip set creates the opportunity for ISVs to create one application port and to run it across two server product lines that can appeal to diverse customers. The IBM team calls this a "two platforms for one port" opportunity.

Common Application Viewpoint

This two-for-one port opportunity exists for most of the applications running on Linux systems today. At the user level of the system, most applications can use the same binary. This reality is reflected in IBM's Software Group statement of support (see Appendix A) for its products on Linux on iSeries and pSeries. For the supported Linux releases from Red Hat, Inc. and SuSE GmbH based on UnitedLinux, there is a single set of CDs common to both iSeries and pSeries systems. There is a single common *glibc* for both series' applications. *glibc* is, of course, the library most used by Linux applications.


Even applications that involve modifying or adding parts to the kernel may be binary compatible for iSeries and pSeries systems. The same code in each distribution above the lowest level of the kernel runs the same way on the two series. The few low-level Linux kernel differences that exist affect the loading of the page tables and the way interrupts are serviced, but the great majority of the code can be the same.

Another great advantage of the POWER architecture is that both 32-bit and 64-bit addressing are fully supported by the same 64-bit Linux kernel on pSeries and iSeries systems. There should be no performance penalty for using either addressing mode since the POWER chips themselves support the two addressing types directly. Because of the additional support for larger scale multi-processor systems in the 64-bit 2.4 Linux kernel, this is the kernel that should be used across the new pSeries and iSeries servers by ISVs who are porting their applications. Using the 64-bit kernel also helps ISVs stage their solution upgrades. ISVs can rapidly move existing 32-bit applications to POWER and

then move up in scalability to a 64-bit space at a planned pace, working with individual pieces of a large solution if that makes sense.

The Linux Environment on POWER Technology

Both iSeries and pSeries can run Linux in a logical partition (LPAR). (pSeries systems can also run Linux in a stand-alone mode.) A partition is a distinct set of hardware resources including processor, memory and I/O. A partition's resources are managed at a partition management console so that resources can be assigned and reassigned to different partitions on the hardware system over a period of time. The ability to change hardware resource allocations gives customers a solution cost efficiency because they can apply their hardware resources where the need is without having separate servers sitting idle part of the time. Unlike separate servers, the idle resource can be applied to another partition while it is not needed for the original workload. The diagram below shows processors on an iSeries being moved dynamically from one workload to another on a scheduled basis.



In addition to Linux partitions, both iSeries and pSeries can run their traditional operating systems in one or more other partitions. A large enough pSeries system can be configured to mix and match up to thirty-two AIX and Linux partitions. In previous and current releases of OS/400, an iSeries system will have at least one OS/400 partition and can also have up to thirty-one other Linux partitions on a large enough system. The ability of these hardware platforms to mix and match Linux partitions with other environments can help ISVs cooperate with other ISVs to create solutions that are more complete than either would deliver alone. Integration functions delivered by the traditional OSs facilitate these cooperative solutions. For example, Samba file sharing is easily set up across all of the partition types. AIX extensions can help administrators create shared user definitions for AIX and Linux partitions. The OS/400 ODBC driver for Linux helps Linux application providers take advantage of existing iSeries customer data residing in DB2 UDB on OS/400.

Up to this point, we have concentrated the discussion on Linux running on POWER hardware. We should also note that Linux applications can be recompiled to use Linux affinity on AIX 5L. These applications also have a high level of binary compatibility with an iSeries environment. The OS/400 compatibility is provided by an AIX runtime, the OS/400 Portable Application Solutions Environment (OS/400 PASE). In this operating mode, the same instruction stream that runs on AIX uses the same libraries as on AIX and gets its kernel services from OS/400. Because of the shared POWER chip set, no emulation is used.

Evaluating the Two-for-One Port Opportunity

While a two-for-one port is the general rule for applications coming to Linux on POWER servers, there are a few application types that are sensitive to the specifics of the system environment. These are applications that have dependencies on specific devices or are very highly tuned to the processor and LPAR environment.

In the area of I/O devices, both iSeries and pSeries web sites contain lists of currently supported devices. The V5R2 iSeries list can be located by going to the iSeries InfoCenter at location

<http://publib.boulder.ibm.com/series/v5r2/ic2924/info/rzalm/rzalmioa.htm>

For pSeries, go to ibm.com/servers/eserver/pseries/hardware/linux_facts.pdf. The devices listed in these two locations have been tested by IBM, and the lists undergo expansion as new devices complete their testing cycles. Also within the Red Hat and United Linux base distributions are the device drivers for the iSeries “virtual devices.” These are means of allowing Linux partitions to share OS/400-owned devices. For virtual disk, OS/400 allocates large chunks of storage that can be used as disks by Linux partitions. These allow, for example, OS/400 to assist with Linux backup and disaster recovery easily.

Part 2 - Getting the Job Done

After an ISV decides to take advantage of the two-for-one port opportunity that POWER technology offers, he or she wants to develop and support the new products efficiently. This is where a platform selection strategy comes in. An ISV is likely to select either iSeries or pSeries as the main POWER development and support platform. The target system selection can be influenced by several factors. One factor is other development platforms that are used in the ISV’s solution repertoire. If more than one development environment is supported on the same POWER footprint, that is a potential efficiency for a development system. For example, ISVs that support both AIX and Linux applications would be likely to choose a pSeries development system. Another factor can be the replication of the most typical customer environment for the new product. For example, if more iSeries customers are early purchasers of the new solution, iSeries may be a good choice for the main development platform.

ISVs that need to maintain a number of development system images would benefit from the iSeries virtualization of Linux storage mentioned earlier or the current iSeries ability to dynamically move processor resources around without rebooting OS/400 or Linux.

Since both iSeries and pSeries platforms support logical partitioning the work of the support team, the test team and the development team of an ISV can be isolated from each other on the same consolidated hardware footprint. Consolidation can notably increase the efficiency of operating the development/test/support environments.

After selecting the main development system, the ISV can do porting work, if some is necessary, on that platform. Generally, only a recompile is needed if the application is running on Linux on another chip architecture. The POWER Linux team has provided several white papers to aid people with porting specifics if there are architecture dependencies involved. More detailed white papers on this topic are located at ibm.com/servers/enable/linux/list.html.

After the port is done on the main development system, it is always advisable to do at least a sanity check on the other platform to be sure that the ISV understands the customers' mindset and environment and that the application works as intended.

To facilitate this kind of activity, PartnerWorld for Developers and eServer Solutions Enablement have implemented programs to help ISVs access hardware in Solution Partnership Centers, as well as providing Remote Development Programs and Test Drives across IBM server platforms. More information on these programs is listed in the references section.

Joe Seiley of AcuCORP, Inc. says this about his two-for-one port activity and the iSeries and pSeries Test Drive engagements his team used.

“We requested an iSeries Linux Test Drive for the original iSeries port of our *extend(tm)* family of legacy application extension solutions. The porting process went smoothly, requiring us to simply recompile and run our standard QA tests. When we moved to the pSeries, all we had to do was take our iSeries Linux solution, install it on pSeries Linux, and once again run our standard QA tests. The results were as successful as the iSeries port. No recompilation necessary.”

Additional POWER Solution Tailoring

For the high performance or highly tuned environment, a significant payback can often be found in using optimizing compilers for POWER that are provided by IBM. The C/C++ compiler is described in a white paper from the POWER Linux team “How to Use IBM VisualAge C++ V6.0 for Linux on pSeries and iSeries” located on the team's website at ibm.com/servers/enable/linux/list.html Binaries produced by the IBM compilers are also compatible between iSeries and pSeries of course.

Every application needs an installation process. From the Linux perspective, the most important aspect of the environment for application installation is the distribution into which the application is being installed. Some file locations will vary from one distribution to another. The same distribution levels that support POWER technology are listed on the iSeries ibm.com/servers/eserver/series/linux and pSeries ibm.com/servers/eserver/pseries/linux web sites. After that, there may be some information about the hardware environment that an ISV should understand to be able to

explain to customers how to install the application onto a system with logical partitions which could be pSeries or iSeries, or onto a stand-alone Linux image on pSeries. Some additional setup and some additional opportunities for application resource sharing may be available in the LPAR environments that are specific to the individual server lines. A high level comparison of Linux on the eServer platforms including iSeries and pSeries can be found in a new Red Book [Linux Handbook: A Guide to IBM Linux Solutions and Resources](#) reference number SG24-7000-00. It can be located by searching the Red Books site www.redbooks.ibm.com.

Summary

In summary, many Linux applications that are being ported to the iSeries and pSeries servers on the POWER architecture require only a single port effort. Most often this effort is only a recompilation. The commonality of hardware across iSeries and pSeries servers and the flexibility of logical partitioning helps ISVs gain maximum benefit from a new chip architecture. These factors also help ISVs appeal to a wide customer base -- small to large business customer environments as well as high-performance customer environments. The commonality also helps maximize the ISV's efficiency in creating the new solution.

Appendix A: Joint Server Group - Software Group Position on Linux iSeries-pSeries Compatibility

- 1) ISVs can expect to see full compatibility between Linux on iSeries and Linux on pSeries with the following minor exceptions:
 - a) I/O – incompatibilities which arise out of the different I/O devices which are supported by each of the eSG [IBM eServer] brands
 - b) Timing – incompatibilities which arise out of processor speed differences and LPAR configurations
- 2) Applications which are either invariant to I/O device (e.g., those which interface exclusively with the file system) or insensitive to processor speed considerations (e.g., those for which performance is not a critical differentiator) will find the two systems fully binary compatible. Applications which are either sensitive to I/O device characteristics or are dependent on timing characteristics of the processor for performance will need to plan for extra testing to mitigate risk that their applications break because of the minor incompatibilities between the two eSG brands.

References

General Linux on eServer Information

IBM eServer Solution Network for Linux ibm.com/servers/enable/linux/

Linux for IBM eServer ibm.com/servers/eserver/linux/home.html

Linux Handbook: A Guide to IBM Linux Solutions and Resources reference number SG24-7000-00 at www.redbooks.ibm.com

iSeries Linux pages ibm.com/servers/eserver/series/linux

pSeries Linux pages ibm.com/servers/eserver/pseries/linux

Porting Information and Programs

IBM eServer Enablement Team white papers ibm.com/servers/enable/linux/list.html

IBM eServer Linux Test Drives ibm.com/servers/enable/linux/testdrive/

POWER hardware references for Linux on iSeries and pSeries

Processors that support Linux on iSeries

ibm.com/servers/eserver/series/linux/hardware.html

I/O devices for iSeries that can be attached to Linux partitions

<http://publib.boulder.ibm.com/series/v5r2/ic2924/info/rzalm/rzalmioa.htm>

PSeries Models and I/O devices that can be attached to Linux partitions

ibm.com/servers/eserver/pseries/hardware/linux_facts.pdf

Acknowledgements

The author is grateful to Steve Dibbell, Chakarat Skawratananond, Matt Davis, Brent Baude and the other Linux solutions Enablement team members for their reviews.

About the Author

Kay Tate is a senior technical consultant in the IBM eServer Solutions Enablement POWER Linux team. She has a Masters Degree in Electrical Engineering from the University of New Mexico. Before coming to IBM, she spent twenty years in the IT industry working primarily on supercomputers, UNIX systems and standards. She joined IBM thirteen years ago in Rochester. After ten years in the IBM doing AS/400 development and strategy, she joined PartnerWorld for iSeries. Since then, she has been working with ISVs to bring their UNIX, Linux and other platform applications to AS/400 and iSeries.