


SPONSORED BY



GEEK GUIDE



Linux on Power

Table of Contents

About the Sponsor	4
Introduction	5
Overview: What Is Linux on Power?	6
Power Architecture	6
Benefits	8
The Power of Virtualization.....	8
Security and Compliance.....	11
High Availability and DR	12
Open Ecosystem	13
Linux on Power Community.....	14
IBM Power Community	15
Conferences	16
OpenPOWER Summit 2016.....	17
LinuxCon 2016.....	17
IBMTechU 2016.....	17
IBM Insight 2016.....	18
ShiftIT	18
Leverage the Power with These Apps	18
Skybot Job Scheduler	19
Halcyon Network Server Suite	21
Conclusion	22

TED SCHMIDT is the Senior Project Manager and Product Owner of Digital Products for a consumer products development company. Ted has worked in Project and Product Management since before the agile movement began in 2001. He has managed project and product delivery for consumer goods, medical devices, electronics and telecommunication manufacturers for more than 20 years. When he is not immersed in product development, Ted writes novels and runs a small graphic design practice at <http://floatingOrange.com>. Ted has spoken at PMI conferences, and he blogs at <http://floatingOrangeDesign.Tumblr.com>.

GEEK GUIDES:

Mission-critical information for the most technical people on the planet.

Copyright Statement

© 2016 *Linux Journal*. All rights reserved.

This site/publication contains materials that have been created, developed or commissioned by, and published with the permission of, *Linux Journal* (the “Materials”), and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of *Linux Journal* or its Web site sponsors. In no event shall *Linux Journal* or its sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

No part of the Materials (including but not limited to the text, images, audio and/or video) may be copied, reproduced, republished, uploaded, posted, transmitted or distributed in any way, in whole or in part, except as permitted under Sections 107 & 108 of the 1976 United States Copyright Act, without the express written consent of the publisher. One copy may be downloaded for your personal, noncommercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Linux Journal and the *Linux Journal* logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners. If you have any questions about these terms, or if you would like information about licensing materials from *Linux Journal*, please contact us via e-mail at info@linuxjournal.com.

About the Sponsor

HelpSystems

HelpSystems empowers IT professionals to excel like never before. Every day more than 9,000 organizations across the globe rely on HelpSystems to automate and simplify system and network management, secure data, and give people simple access to information they need. For critical needs like IT and business process automation, system security, network mapping, document management, and business intelligence, HelpSystems makes IT lives easier and keeps businesses running smoothly. Learn more at <http://www.helpsystems.com>.

Linux on Power

TED SCHMIDT

Introduction

Until recently, IBM's Power Platform was looked upon as being the system that hosted IBM's flavor of UNIX and proprietary operating system called IBM i. These servers often are found in medium-size businesses running ERP, CRM and financials for on-premise customers. By enabling the Power platform to run the Linux OS, IBM now has positioned Power to be the platform of choice for those already running Linux that are facing scalability issues, especially customers looking at analytics, big data or cloud computing.

Running Linux on IBM's Power hardware offers some obvious benefits, including improved processing speed and memory bandwidth, inherent security, and simpler deployment and management. But if you look beyond the impressive architecture, you'll also find an open ecosystem that has given rise to a strong, innovative community, as well as an inventory of system and network management applications that really help leverage the benefits offered by running Linux on Power.

Overview: What Is Linux on Power?

Power Architecture: Power is an acronym for Performance Optimization With Enhanced RISC, which originally was intended as the architecture for mid-range servers and workstations. A Reduced Instruction Set Computer (RISC), of course, improves performance tremendously, allowing you to move system functionality from hardware to software. Because the Power architecture increases performance of parallel processing by increasing addressable memory (16TB) and providing eight threads for every CPU, for a total of 96 concurrent threads, there are incredible improvements in the time it takes to access memory. Faster queries and reduced time for data to be brought into memory is critical for applications like big data.

Compare these to stats for the x86 processor, which can offer only 30 concurrent threads at a maximum processor speed of 3.4GHz, on-chip cache of 37.5MB, and only 1.5TB of addressable memory per socket (Table 1), and there's just no contest. Power clearly is the platform for, well, for power.

Table 1. IBM Power8 vs. Intel

Spec	IBM Power8	Intel Haswell
Frequency	3.52GHz	2.3GHz
Cores	12	16
Logical Cores	96 SMT	32 HT
Cache	96MB	40MB
Memory Bandwidth	384GBps	68GBps
Max GFLOP/s (single)	676	1,178
Max GFLOP/s (double)	338	589

But beyond impressive improvements in processing speeds and concurrent threads, IBM also added something called Coherent Accelerator Processor Interface (CAPI), which allows hardware to bypass the operating system and communicate directly with the CPU. This results in even faster data access speeds and more data being accessed—another big boost for data analytics.

Finally, Power is BI-endian, meaning it supports both big-endian (BE) and little-endian (LE) data management. Although RISC defaults to big-endian—that is, treating the number one hundred twenty-three as “123”—it supports little-endian, treating the same number as “321”.

Intel x86 processors support little-endian Linux, which means that any application currently running on Linux on an Intel x86 processor also can run on Power

on Linux. Because Power processors support little-endian Linux, customers can realize big improvements in processing power by moving from x86 processors to Power with relative ease and no additional development.

To prove the point, Canonical (Ubuntu Linux distributor) moved 40,000 Linux applications over to the Power platform in only 160 days. As well, IBM already has more than 80% of its own Linux applications ported over to the Power platform, and SUSE and Red Hat have Linux releases running on Power. Anyone currently running Linux on an x86 architecture looking for more power and speed to support analytics or cloud computing, as well as commitment from IBM and the Linux community, will take notice of that.

Benefits

It's clear that the RISC-based architecture of IBM Power provides a much faster and more powerful platform than x86. And it has been proven how easy it is to migrate to Linux on Power for current users of x86 processors running Linux. So, let's look at some of the benefits, other than processing speed and ease of migration, that Power on Linux offers.

The Power of Virtualization: Power servers are built specifically to run multiple virtual machines (VMs) without increasing latency by increasing priority and avoiding interference of latency-sensitive tasks. IBM offers several options for virtualizing machines on a Power system, although VMware is not currently among them.

The first option, PowerVM, will support running

GEEK GUIDE ► LINUX ON POWER

multiple LPARs of different versions of UNIX, including AIX, Ubuntu, Red Hat and SUSE at the same time and still support IBM i for your core business application. A variation of this option, called “PowerVM for PowerLinux” supports only Linux VMs. As an alternative to PowerVM, PowerKVM is the open-source-based virtualization solution offered on Power8 servers (versus the previous generation of Power7). IBM provides the ability to switch between PowerVM and PowerKVM with a simple reboot of your VM. Table 2 compares the two options on a Power8 platform.

Table 2. Terminology on PowerVM and PowerKVM

PowerVM	PowerKVM
FSP	Flexible Service Provider (FSP)
Hypervisor	Host, Hypervisor
Logical Partition (LPAR)	Virtual Machine (VM), Guest, Virtual Instance
PowerVM hypervisor driver (pHyp) firmware	Open Power Abstraction Layer (OPAL)
Virtual I/O server (VIOS) and the PowerVM hypervisor driver	KVM Host Userspace (QEMU)
Integrated Virtualization Manager (IVM) with Hardware Management Console (HMC)	Host userspace tools based on the libvirt API, including virsh
Integrated Virtualization Manager (IVM with Hardware Management Console (HMC)	Management of virtualization through a browser-based user interface with Kimchi. Kimchi utilizes libvirt to manage KVM.
Integrated virtualization manager (IVM) with Hardware Management Console (HMC)	Command-line message-based hardware management interface to manage IPMI-enabled devices on remote host with IPMtool.

The benefits of using PowerKVM include memory virtualization, I/O virtualization and live VM migration. Memory virtualization allows memory to be over-committed and swapped out when underutilized, just like in a regular Linux system. I/O virtualization offers a wide range of I/O resource sharing options, including PCI pass-through to allow VMs to access dedicated devices directly, and I/O configurations that include Fiber Channel, NFS, iSCSI and Ethernet. If you do decide to leverage the benefits of KVM, be aware that PowerKVM requires the Open Power Abstraction Layer (OPAL) interface to access the hardware, and because of that, it runs only on Power8 servers. Be sure to evaluate all the configuration options carefully before making a final choice.

PowerKVM also provides Red Hat Enterprise Virtualization (RHEV) and bare-metal virtualization. The former option is important because Red Hat does not yet support RHEL if it is running on PowerKVM, so you need to use its virtualization. The latter option, which is really no virtualization, is important for those looking for the performance boost of 3–20% that comes from removing the virtualization layer from between the hardware and the virtual machine. Only Ubuntu is enabled for bare metal at this time.

So, PowerVM provides the ability to create multiple LPARs of different operating systems, with the exception of the PowerLinux variation, which supports only Linux. PowerKVM is the Power8 solution that takes the next step by providing memory and I/O virtualization, along with live VM migration. These are important to anyone looking at cloud computing, but be aware of some of the limitations when choosing which virtualization option to use.

PowerSC also ensures the segregation of virtual networks, which ensures that security compliance is maintained, and provides local firewall services on the virtualization layer, which helps improve performance and reduce network resource consumption.

Security and Compliance: Because the hypervisor for Power resides directly in the firmware of the system, it makes the virtual machines just as secure as the physical servers on which they reside. In fact, the Department of Homeland Security's United States Computer Emergency Readiness Team (US-CERT) and MITRE Corporation reported zero vulnerabilities against PowerVM. Both organizations are global leaders in establishing and auditing cyber-security protocols. As well, HIPAA, DOD, PCI, COBIT and other compliance protocols are preconfigured into PowerSC, the security and compliance component of Power Systems, and can be set up to notify an administrator automatically if a VM falls out of compliance. Through compliance reporting, PowerSC even offers suggestions for bringing non-compliant VMs back into compliance. PowerSC also ensures the segregation of virtual networks, which ensures that security compliance is maintained, and provides local firewall services on the virtualization layer, which helps improve performance and reduce network resource consumption.

High Availability and DR: Power on Linux leverages native Distributed Replicated Block Device (DRBD) to provide HA and DR capabilities. Using RAID-1, DRBD mirrors an entire block device to another networked host during runtime, which allows the development of high-availability clusters for block data.

DRBD supports a variation on synchronous and asynchronous data replication that ends up providing more data protection than asynchronous at a small cost in performance. To ensure data integrity, DRBD uses the Linux Crypto API and calculates message integrity codes to go with data moving between nodes. This allows the receiving node to validate its incoming data and request retransmission when an error is found. Additionally, DRBD permits the verification of local and peer devices in an on-line fashion (while input/output occurs). Finally, DRBD provides several automatic methods for recovering from split-brain errors, when each node thinks it is primary.

When thinking about providing high availability (HA) for PowerKVM on Linux, SUSE also provides an excellent option with its Linux Enterprise High Availability Extension. It comes fully integrated as an extension with SUSE Linux Enterprise Server, which avoids issues that come from using third-party products. It provides a great 24x7 HA clustering solution that includes IP load balancing, TCP disk replication and shared and scaled data access by multiple nodes via the Oracle Cluster File System.

In summary, IBM's Power platform provides multiple performance benefits over Intel's x86 system. The RISC architecture provides impressive speed and parallel

processing advantages over the x86 system. Combined with memory availability and a CAPI that provides the CPU with direct access to external flash storage, this is a big boost for anyone doing big data analytics. Power also supports BE and LE, which means migrating Linux applications from x86 is painless and quick. Finally, Power offers significant advantages in virtualization, security and HA/DR, which serve the needs of cloud computing very well.

Open Ecosystem

In 2013, IBM founded the Open Power Consortium (now called the OpenPOWER Foundation) with Google, Mellanox, NVIDIA and Tyan in an effort to create an entirely open ecosystem around the Power architecture. The goal is to drive innovation by making all the intellectual property around the Power architecture completely available to all OpenPOWER Foundation members. From Google's investment in building its own hyperscale processors based on Power8, to NVIDIA's work on providing an alternative to IBM's CAPI interface for access to Graphics Processing Units (GPUs), to Tyan's customer reference system (based on the Power architecture, which allows end users to deploy software that is customized to their unique requirements), IBM's willingness to create an open ecosystem around the Power architecture appears to be working, with 157 members in 22 countries.

But to keep things moving in the right direction, IBM continues to work closely with ISVs (Independent Software Vendors), like HelpSystems who already have investments in

Creating and working with a broader community helps make Power more relevant through open innovation, and the OpenPOWER Foundation helps promote collaboration on multiple levels, from chip architecture, to hardware acceleration, to systems management software.

Linux, AIX and IBM i solutions. These partners know what customers want and are looking at ways to help IBM expand and manage the capabilities offered by running Linux on the Power platform. They appreciate the value of big data and cloud computing, and they understand the value that Power on Linux offers their customers.

Linux on Power Community: Creating and working with a broader community helps make Power more relevant through open innovation, and the OpenPOWER Foundation helps promote collaboration on multiple levels, from chip architecture, to hardware acceleration, to systems management software. The Foundation already has 11 chartered workgroups, such as Systems Software, Memory, Accelerator and Hardware Architecture, just to name a few. These workgroups are focused on delivery of projects like LE Linux KVM, OpenCL SDK to OpenPOWER Platform interface and IO Device Architecture (IODA2), all of which will increase the platform's accessibility to

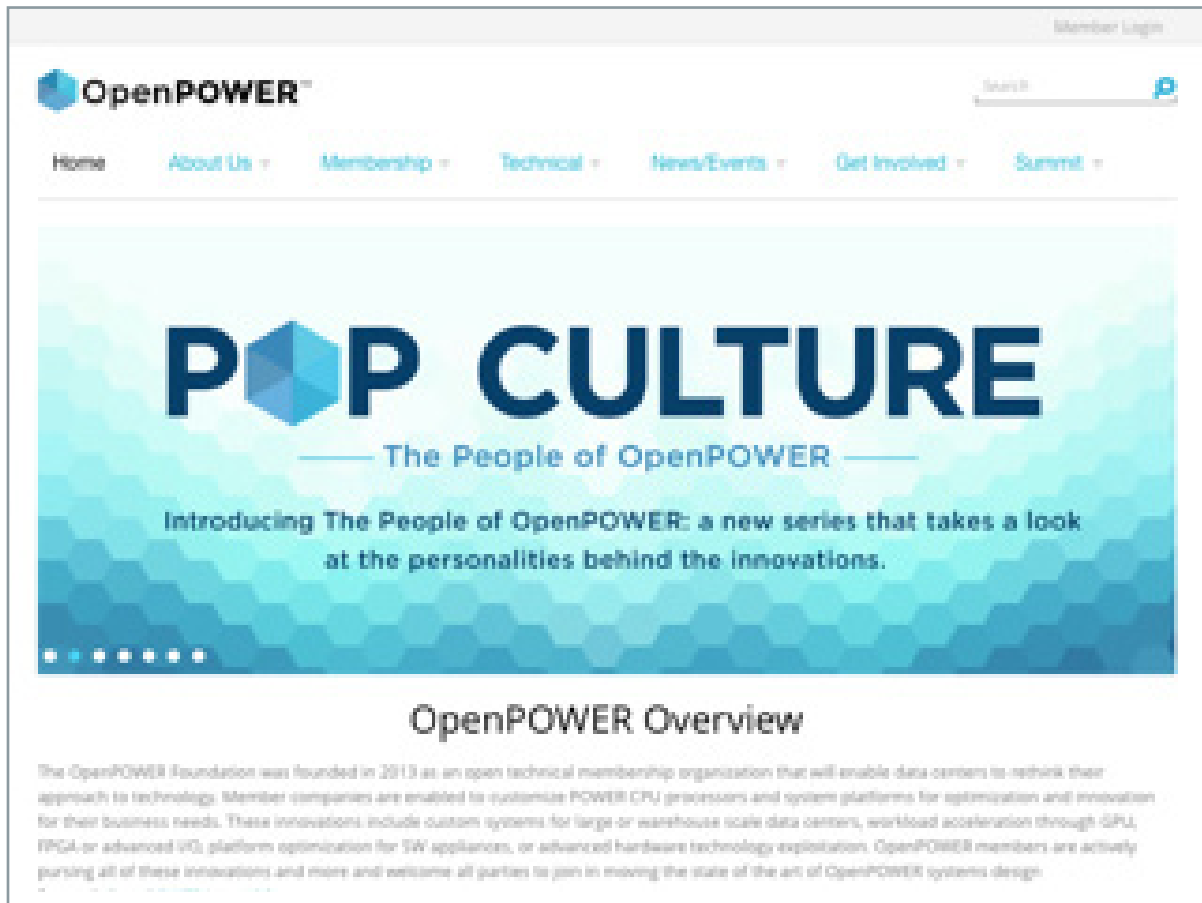


FIGURE 1. OpenPOWER Foundation:
<http://openpowerfoundation.org>

channel partners. IBM and the OpenPOWER Foundation are committed to expanding the open-source ecosystem for Linux on Power.

IBM Power Community: In addition to the OpenPOWER Foundation, IBM has a strong on-line community that is geared more toward individuals in the Linux on Power community, where the Foundation is geared more toward channel partners. IBM's on-line community provides an array of technical resources for Linux software on Power systems,



FIGURE 2. The Linux on Power Community:
<https://www.ibm.com/developerworks/group/tpl>

including direct access to other Linux on Power users and individuals on the development teams for both Linux and IBM Power technologies. From blogs that herald releases of new features, to wikis that provide access to detailed technical knowledge, to open forums where you can find access to the community's top experts in Linux on Power, the IBM Power Community offers something for every need.

Conferences

In addition to collaborating on-line with others in the Linux on Power Community, it's obviously important to get some hands-on experience and training, and to establish working relationships with other professionals.

OpenPOWER Summit 2016: The OpenPOWER Foundation held the 2016 OpenPOWER Summit in early April, building upon the great success of the Foundation's inaugural Summit in 2015. This event featured speakers and demonstrations from the OpenPOWER ecosystem, including industry leaders and academia sharing their technical solutions and state-of-the-art advancements.

For more information on the OpenPOWER Summit, see <http://openpowerfoundation.org/openpower-summit-2016>.

LinuxCon 2016: The Linux Foundation offers events that are collection points for maintainers and developers across multiple important open-source software projects in the enterprise, networking, embedded and cloud infrastructure. Any company (or person) that wants to leverage the power of open source, especially Linux on Power, should strongly consider attending LinuxCon.

If you rely on, or are even considering Linux on Power, it really makes sense to be a part of Linux Foundation events, because they are a great place to meet and network with the people who are actively involved in the Linux on Power community.

More information about LinuxCon is available at <http://events.linuxfoundation.org/events/linuxcon-north-america>.

IBMTechU 2016: IBM also offers its own events across the country every year. If you want to grow your Linux on Power knowledge and skills, IBMTechU offers multiple opportunities with in-depth technical training at one of the IBMTechU Comes to You events across the country.

There are hundreds of sessions led by IBM product and industry experts where you can learn about Linux on Power or sharpen existing skills. IBM TechU offers hands-on labs, demos and workshops geared to all skill levels, so it's easy to find whatever you're looking for. IBM TechU also offers the opportunity to test-drive new products from IBM and IBM channel partners in the Solution Center.

You can learn more and register at <http://www-304.ibm.com/services/learning>.

IBM Insight 2016: IBM Insight 2016 is slated for October 2016 and is another way to stay connected with the Linux on Power community with more than 1,600 in-depth educational sessions including technical training, hands-on labs and industry use cases presented by more than 1,000 industry experts. It also includes certification opportunities.

ShiftIT: Channel partners also provide a great wealth of conference opportunities for staying connected to the Linux on Power community. For instance, this year, HelpSystems will be sponsoring ShiftIT, an opportunity for Linux on Power professionals to get equipped with the tools, skills and perspectives they need to leverage and secure the power and performance offered by a Linux on Power data center.

Visit <http://www.helpsystems.com> for more information on ShiftIT.

Leverage the Power with These Apps

Earlier I mentioned existing IBM channel partners

With key features like a central monitoring interface, automated alert notifications and high-availability and auditing capabilities, Skybot quickly helps you reap the performance benefits offered by Linux on Power and makes the argument for moving Linux from Intel to Linux on Power.

who are making great strides in leveraging the open ecosystem of Linux on Power and providing tools to customers to harness the power of this platform. HelpSystems, which started providing automated scheduling solutions for IBM S/38 in 1982, is a great example of an ISV partner that helps strengthen the argument for moving to Linux on Power.

Skybot Job Scheduler: Using a self-contained HTTP server and database that run on Linux on Power and connect to agents installed on all of your managed servers (Linux, UNIX, IBM i and Windows), Skybot offers an easy way to automate your data-center operations. With key features like a central monitoring interface, automated alert notifications and high-availability and auditing capabilities, Skybot quickly helps you reap the performance benefits offered by Linux on Power and makes the argument for moving Linux from Intel to Linux on Power.

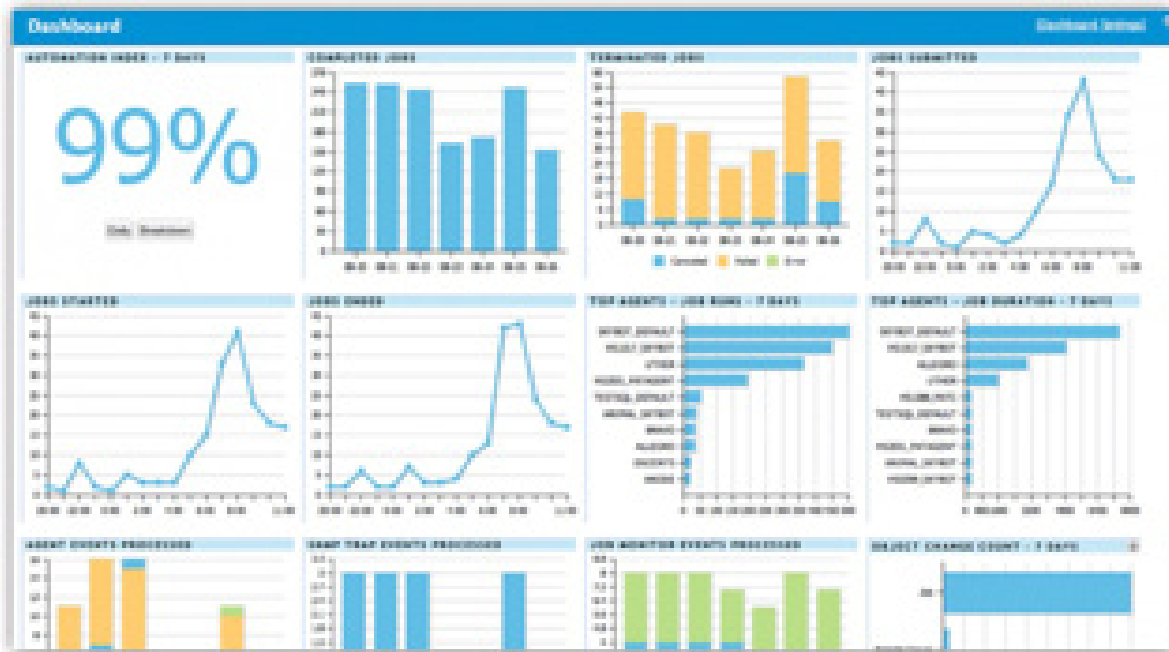


FIGURE 3. HelpSystem’s Skybot

Energy Solutions is an energy utilities consulting company in Oakland, California, that manages multiple Web sites with incentive programs and also provides public portals with energy-use reporting for the CA utilities commission. The company accomplishes all this with less than ten servers, two transaction processing systems and Skybot Scheduler, the enterprise job scheduling and workload automation software.

Before Skybot, Energy Solutions had separate schedulers across multiple servers, which required labor-intensive maintenance and manual intervention to check on whether batch jobs were complete. Development and maintenance of custom scripting required a lot of expensive developer time. Time-based

Halcyon's monitoring software pays close attention to critical IT systems and core applications on which business depends.

jobs were inefficient, redundant and produced job logs that provided little or no value.

By installing Skybot, Energy Solutions was able to create a centralized view of its entire enterprise infrastructure. The company replaced manual job checks with a consolidated job history and implemented automatic notification of any delays or errors. It also replaced custom-written scripts with flexible, cross-system scheduling in Skybot Scheduler, and it finally was able to create sensible job schedules that included real-time triggers based on system events like file arrivals.

Halcyon Network Server Suite: HelpSystems also offers the Halcyon Network Server Suite for Linux on Power. This suite provides Linux, Windows, AIX and VIOS monitoring alongside applications that are hosted on Linux on Power to ensure maximum availability.

Halcyon's monitoring software pays close attention to critical IT systems and core applications on which business depends. Early warnings of any potential problems or "bottlenecks" are sent to administrators before they affect the business. Through automation, Halcyon also allows IT teams to create self-corrective actions to problems that arise, helping to prevent

recurrence of problems in the future. Normal tasks, checks or jobs, which typically might be done manually, also can be automated, freeing up additional development resources, which helps reduce costs and allows IT teams to focus on other initiatives to drive the business forward.

Halcyon also comes with a cloud-based reporting solution that shows management that the IT team is achieving the service levels expected by the business. The same reporting solution can be used to flag any issues as they arise so they can be resolved quickly.

Halcyon helps companies prevent business disruption, increase the efficiency of their IT services and save costs. On average, Halcyon's solutions increase IT efficiency by 30–40%. For most companies, this equates to significant savings and an impressive return on investment for Linux on Power shops. And, HelpSystems is just one of the many IBM Linux on Power ISV partners who is participating in the collaborative opportunities offered by IBM's Linux on Power open ecosystem.

Visit <http://www.helpsystems.com> for more information on both of these IBM Linux on Power channel partner tools for getting the most out of Linux on Power.

Conclusion

IBM has long provided top performance with its Power Systems servers. The RISC-based Power platform far surpasses Intel's x86 in CPU speed and parallel processing and offers more memory availability and a CAPI that provides direct CPU access to external flash

storage. For big data analytics, this is requisite. Power supports big-endian and little-endian data formatting, which allows existing x86 customers to migrate their Linux applications quickly and easily. And, for any enterprise interested in cloud computing, like Google, Power offers impressive advantages in virtualization, including support for LPARs running different operating systems, and virtualization of memory and I/O. With security that's built in to the hardware, high-availability and disaster recovery, Power also provides the answer for those with compliance and audit concerns.

With the more recent move to an open ecosystem and the support of strong channel partners like HelpSystems, IBM has positioned Linux on Power to be the enterprise-computing platform of choice, especially where the demands of data analytics and cloud virtualization are necessary. Not since AIM was dissolved has there been as strong of an open ecosystem as there is around Linux on Power. By building a strong community around a technology that offers scalable, secure virtualization, multicore options for any workload, the ability to connect the CPU directly to external hardware and ease of migration from x86 hardware, IBM has positioned Linux on Power as the enterprise platform of the future. ■