

Helping Solve Nature's Mysteries ▪ 46

Iowa State University turns to Linux to research genomes. *By Shirley S. Savage*

A Seamless Enterprise ▪ 50

Linux is a good fit with Success Apparel's business goals. *By James Mathewson*

Past, Present and Future ▪ 54

Three executives discuss the whys and wherefores of IBM's approach to open source beyond Linux. *By Jim Utsler*

Creating a Community ▪ 59

Palm Beach Community College learned that consolidation and virtualization can result in more than just lower costs. *By Jim Utsler*

Departments

Executive Viewpoint ▪ 8

Sharing in the Success of Open Source

Industry Watch ▪ 10

IBM's Migration Factory Eases Transitions

Trends ▪ 12

Linux Hits the Mainstream

Linux Line ▪ 63

Customers Discuss Benefits of Open Source

Parting Thoughts ▪ 64

Conquering Open-Source Fears

Page 54

Page 50

Page 59



There's a new guide for systems administrators. Download a free technical reference for migration from Solaris to Linux technology-based systems at www.ibm.com/linux.

LINUX

Executive Report from IBM

EDITORIAL

FOUNDER, EXECUTIVE EDITOR & PUBLISHER

Barry R. Mittelman
Senior Marketing Manager, Competitive Marketing Lead
Worldwide Linux & Open Source, IBM Corporation
(bmittel@us.ibm.com)

CUSTOM PUBLISHER: MSP TECHMEDIA

ASSOCIATE EDITORS

Jeff Cork
Evelyn Hoover
Caroline Joyce
Doug Rock

WRITERS

Aaron Dalton
Elliot King
Kristin Lewotsky
James Mathewson
Joe McKendrick
Maureen Nevin
Shirley S. Savage
Neil Tardy
Jim Utsler

PRODUCTION MANAGER

Jonathan Benson

ART DIRECTOR

Kim Dalros

DIRECTOR OF SALES

Mari Adamson-Bray

CIRCULATION MANAGER

Linda Holm

© Copyright 2006 by International Business Machines (IBM) Corporation. This magazine could contain technical inaccuracies or typographical errors. Also, illustrations contained herein may show prototype equipment. Your system configuration may differ slightly. This magazine contains small programs that are furnished by IBM as simple examples to provide an illustration. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. All programs contained herein are provided to you "AS IS." IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE EXPRESSLY DISCLAIMED. All rights reserved.

All customer examples cited represent the results achieved by some customers who use IBM products. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions. Information concerning non-IBM products was obtained from the products' suppliers. Questions on their capabilities should be addressed with the suppliers.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.

The following are trademarks (marked with an ®) of the International Business Machines Corporation in the United States and/or other countries. A complete list of IBM Trademarks is available online (www.ibm.com/legal/copytrade.shtml).

AIX	IBM logo	S/390	TotalStorage
BladeCenter	MVS	System i	VM/ESA
CICS	OpenPower	System p	VSE/ESA
DB2	OS/390	System Storage	WebSphere
ESCON	OS/400	System x	z/OS
FICON	POWER	System z	z/VM
IBM	RS/6000	Tivoli	

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc., in the United States and other countries. Microsoft, Windows and Windows NT are trademarks of Microsoft Corp., in the United States and other countries. Intel, MMX and Pentium are trademarks of Intel Corp. in the United States and other countries. UNIX is a registered trademark of The Open Group in the United States and other countries. SET is a trademark owned by SET Secure Electronic Transaction LLC. Linux is a trademark of Linus Torvalds in the United States, other countries, or both. Other company, product and service names may be trademarks or service marks of others.

Reprints: To order reprints, contact Kelly Carver at (612) 336-9280.

The *Linux Executive Report* from IBM is published by the International Business Machines Corporation, Route 100, Somers, NY 10589. The articles in this magazine represent the views of the authors and are not necessarily those of IBM.

Publications Agreement No. 40916539, Canadian Return Address, World Distribution Services, Station A, PO Box 54, Windsor Ontario Canada N9A 6J5.

Printed in the U.S.A.

Sharing in the Success of Open Source

What makes Linux* so influential? Is it the fact that it's sent a considerable shock wave through the computer industry? Is it because it's helped facilitate the collaboration of programmers, engineers, scientists and organizations across the planet? Does it have anything to do with helping less-developed nations facilitate access into the global marketplace?

All of these facts contribute to the positive influence of Linux. However, the greatest influence of Linux may be that it validated the powerful philosophy and business model of using open-source technology to more efficiently deliver a shared, industry-wide open standard to customers. Because of the market success Linux has achieved, scores of other technologies that are "open source beyond Linux" have credibility and success as a way to implement or accelerate other shared industry-wide open standards. The same repeatable model has fueled the success of technologies such as Eclipse, several significant Apache projects and the Aperi project (read more about these topics beginning on page 54).

In fact, companies that innovate on top of open standards are advantaged because resources are freed up for higher-value work and because market opportunities expand as the standards proliferate. And customers win because vendors who use open source to more efficiently deliver implementations of those open standards can use the freed-up resources to deliver more value-add to customers on top of the standard. Customers like open standards, as they want simplicity and common standards help to simplify multivendor IT environments. Everyone wins in an environment that uses open source to share implementations of open-technology standards. The success of Linux has been a critical catalyst in making that idea a reality.

All of the Linux and open source-related attributes listed above have measurable business benefits. One attribute that shouldn't be overlooked, however, is flexibility. Companies

that have adopted Linux as their operating system of choice are often doing so because of its flexibility. With Linux, an organization's product development cycle is no longer shackled to the development schedules of proprietary software companies. Mark VandenBrink, senior director and chief platform architect for Motorola Mobile Devices explains in a quote from this issue's Trends column:

"We went with Linux because it provides a rate of innovation we found compelling," VandenBrink says. "When we looked at Linux, we saw the ability to push things strategically faster along a certain path, without having to wait on, or be gated by, somebody else's release schedules. Linux provides a way to rapidly innovate and get new features into a device with a lot more freedom."

Linux and open-source initiatives provide individuals and organizations a common and simplified technology foundation that allows them the freedom and flexibility to build the solutions that best serve their business needs. You'll find examples of companies that have used this model to realize success in this issue of the *Linux Executive Report*. All the best.



A handwritten signature in black ink that reads "Scott Handy". The signature is fluid and cursive, written in a professional style.

Scott Handy
Vice President Worldwide Linux and Open Source, IBM



Easing Transitions

Inside IBM's Migration Factory: Organizations choose Linux to replace Sun operating systems **BY ELLIOT KING**

ONE CONSTANT IN IT management is change. Even if an application is meeting its service levels today, it may have to be migrated to a new platform tomorrow. "There are a number of reasons why a customer would migrate," says Kevin Galloway, technical services manager for the IBM* Migration Factory, a part of the Application-management Services group within IBM Global Business Services. These reasons could be a particular hardware platform coming to the end of its life, a company looking to consolidate servers due to acquisition and/or reduce IT cost, an enterprise undertaking a periodic technology refresh or simply choosing a new vendor.


Those factors are coming into play as many organizations migrate their infrastructures from Sun Microsystems Solaris operating systems (OSs) to Linux* running on IBM hardware. With some Sun Microsystems computers reaching the end of their useful life, companies are moving to smaller systems with fewer processors that offer the same or higher amount of processing power. "That reduces licensing and operating costs for them," Galloway noted. Moreover, he added, power use and heat generation have emerged as significant challenges in many data-center operations.

IBM's Migration Factory has skills to cover various kinds of workloads. "We break them down into key components: infrastructure (e-mail, networking, Web)," says Galloway.

The Migration Factory's primary objective is to lower risk and drive down the

cost of migrating to an IBM platform. "We focus on understanding the business, technology and strategic aspects of migration. It's not all about speeds and feeds," Galloway said. "There are technology issues, working risk factors to minimize the customers' fears about risk."

Working with the different IBM hardware brands, the Migration Factory helps enterprises understand the resources, time



The Migration Factory's primary objective is to lower risks and drive down the costs of migrating to IBM platforms from original platforms.



and support needed to complete a successful migration project. "We help them understand what goes into a migration project," Galloway said. The IBM Migration Factory team assists with the complete migration process from initial planning to using the migrated technology in production.

With more than 20 years of experience and more than 4,000 successful migrations completed, the Migration Factory has developed a methodology and set of best practices that serve to help ensure a project's success. "One of the core parts of the methodology and one that helps mitigate the risk is the extent to which we introduce testing into the migration process," Galloway said. IBM performs its own testing and encourages the customer to run their own performance testing as part of the acceptance process. "We haven't simply gone in and moved code over and then said to the customer that it's their job to get the application running,"

Galloway said. "When the migrated application fulfills the requirements of the IBM and customer-defined testing, customers see that they have a working application on Linux."

Migrating an application to the Linux platform is a team effort. The Migration Factory connects customers to IBM's educational offerings so enterprises can improve their Linux administration and development skills. "A key aspect of

moving from one platform to another is understanding how to manage, operate and develop on the new platform," Galloway said.

Since Linux is a fully supported OS, IBM offers ongoing support following a migration and provides customers with the benefits of IBM's close relationships with its Linux partners Novell and Red Hat. "It is very much a team IBM effort," Galloway said. "We have expertise in our group and also have access to other people at IBM who are supporting, developing and enhancing the existing open source code."

According to Galloway, the number of migrations from Solaris to Linux is picking up, with telecommunications and financial services industries leading the way. "They want to be at the leading edge and that is driving acceptance," he said. We have experienced an increase in the number of applications available on Linux and increases in the number of migrations to Linux as well, he said.

Smoothing the Transition From Solaris to Linux

IBM has just published a new guide, “Solaris to Linux* Migration: A Guide for System Administrators”

(www.redbooks.ibm.com/abstracts/sg247186.html), which provides an easy-to-follow step-by-step guide for moving from Solaris 8 or 9 to Linux on any type of vendor hardware platform.

From an operational point of view, Solaris and Linux are very similar environments.

However, as with any two operating systems (OSs), there are nuances and distinctions of which administrators at migrating enterprises need to be aware. Some key “pain points” that the guide helps alleviate include the following:

Software-package management:

The Solaris package format is based on the System V Interface Definition for Application Binary Interface (ABI) and has tools for managing these packages. Red Hat and Novell SUSE Linux both use the Red Hat Package Manager (RPM), which includes features for package management such as dependency and signature checking and other advanced options.

Patch management: The patch management process in Solaris requires the installation of the base version package first, followed by all the associated patches. Linux should be a welcomed transition for administrators doing patch management. Instead of creating patches, vendors create full packages of fixed software to simplify package maintenance. If a system administrator needs to install a new package, the process only requires installing the latest version.

Patch activation: To activate a patch in Solaris, the system may have to regenerate its devices during the reboot if the fix is for driver components. In most kernel or library patches, Sun recommends installing in single user mode. This allows the system to

be immediately rebooted after applying the update, while minimizing the potential for a loss of user data. In Linux, however, there is no such thing as reconfiguration reboot or device regeneration. Sometimes it is possible to update kernel modules without rebooting the system by removing the affected module from kernel, updating the kernel module, and reinserting the new kernel module.

Disk management: In both Solaris and Linux, before a disk can be sliced for use, the disk has to be recognized by the operating system. However, Linux is more automatic. Solaris requires setting up a file to recognize specific targets and LUN numbers for new SCSI devices in the system. In Red Hat and SUSE Linux, device files are automatically created every time a system boots, so a reconfiguration reboot is not necessary.

Disk file systems: Solaris natively only supports the old UFS file system for hard disks which has changed in Solaris 10. The other types of disk-based file systems are HFS (for CD-ROM), PCFS (for diskettes), and UDF (for DVD).

Administrators will love the options and tools for modern file systems in Linux. There are several disk file systems from which to choose. On Red Hat, the default disk file system is ext3, and on SUSE, it is reiserfs.

Basic user administration: The Solaris version of LDAP uses the iPlanet Directory Server, while both Red Hat and SUSE Linux use OpenLDAP, a software package founded with support from IBM. This package provides a stand-alone

and update replication daemon, as well as libraries, utilities, Java* classes and sample clients. OpenLDAP is also available for Solaris.

Security hardening: Most Solaris administrators are familiar with Yet Another Solaris Security Package, YASSP and the Solaris Security Toolkit, used to more efficiently harden a Solaris system. Sun also supported Solaris Extensions (as an add on—not integrated). Security hardening is integrated into Linux distributions. One of the most popular tools used to implement system hardening is called Bastille. Bastille Linux is an open-source project aimed at hardening a Linux system by changing some of the system’s settings based on security best practices.

Installation: The good news here is that Solaris and Linux are virtually identical in installation and setup although administrators should find modern Linux setups much easier than older Solaris or Solaris 10. As with a Solaris installation, Linux will perform operations such as probing the hardware for devices and load the appropriate device drivers, assigning a host name, setting up a network setup type (whether DHCP or static IP), and prompting for software bundles or packages to install.

Joe McKendrick is a research consultant specializing in IT and organizational development trends, authoring special reports for IDC and Gartner. Joe can be reached at joe@mckendrickresearch.com.



ONLINE: Go to www.ibm.com/linux to download a free technical reference for migration from Solaris to Linux technology-based systems and to get more information about migrating from Solaris to Linux.

Of course, the Migration Factory handles more than Solaris to Linux migrations. “Linux is becoming a target from Windows*, HP-UX, Tru 64, UNIX, VMS and Tandem as well,” Galloway said. “Workloads are moving from all of

those platforms to Linux.” All kinds of IT users can benefit from the expertise of IBM’s Migration Factory.

Elliot King is a professor at Loyola College in Maryland where he specializes

in new communication technology. Elliot has written five books and several hundred articles about the emergence and use of new computing and communication technology. Elliot can be reached at eking212@comcast.net.

Going Mainstream



Linux is now found in countless systems and devices worldwide

BY JOE MCKENDRICK

Linux* has been successful as a server and client-side operating system (OS) in recent years, but such visible enterprise deployments are only part of the story. Linux is also working behind the scenes, powering countless systems and devices, including network equipment, printers, cell phones, smart mobile phones, kiosks, point-of-sale terminals, digital-media processors, routers, firewalls, storage devices, medical-imaging equipment, industrial-control applications and private-branch exchanges. Linux now runs on a huge array of processors that power these devices, including IBM* PowerPCs, ARM embedded processors, Intel* XScale, Intel Pentium M, AMD Geode, NEC MPCore and countless others.



“When you connect to the Internet through your local ISP, the wireless router you use to connect your notebook is very likely running Linux, as is your telecommunications provider.”

—Bill Weinberg, senior technology analyst, Open Source Development Labs

The embedded-Linux market is now a growing part of the mainstream embedded market. Recent surveys from Venture Development Corporation (VDC), for example, calculate that between 11 and 15 percent of developers now use Linux in their embedded projects—making this a leading OS choice. (By contrast, the second and third leading categories—Wind River VxWorks and Microsoft* Windows* XP Embedded—drew 10 percent and 6 percent, respectively.) Linux use is expected to grow to an almost 20-percent share within the next two years. For 32-bit implementations, the share of embedded Linux is already at 25 percent, VDC found. VDC also put the embedded-Linux market at \$134 million in shipments during the last year, a 46-percent jump from the previous year.

While unaware of it, end users touch embedded Linux every day as they go about their business. “When you make a wireless phone call, you’re likely to traverse two or three copies of carrier-grade Linux,” says Bill Weinberg, senior technology analyst at Open Source Development Labs (OSDL). “Linux may run in your phone and also inhabit your network. When you connect to the Internet through your local ISP, the wireless router you use to connect your notebook is very likely running Linux, as is your telecommunications provider.” An impressive roster of brand-name equipment manufacturers has adopted Linux, including IBM, Sony, Motorola, Phillips, Panasonic and Siemens.

Embracing Linux

Motorola has been increasing its presence in the Linux and open-source space, recently launching a Web site aimed at reaching open-source developers. “What’s key to a consumer electronics company such as ours is the brand and the ability to innovate around that brand,” explains Mark VandenBrink, senior director and chief platform architect for Motorola Mobile Devices. “Linux enables us to innovate in the way we need to. All of our strategic innovation now takes place on the Linux platform.”

The platform providers that power these devices also are embracing Linux. MontaVista Software, started in 1999, is perhaps one of the earliest embedded-platform vendors to build on Linux. Currently, MontaVista, an IBM business partner, supports more than 200 Linux-support packages for evaluation and deployment computer boards from companies

worldwide. Weinberg, an original founder of MontaVista, related that he had “spent many years with IBM in helping to embed Linux on PowerPC and other processors, where Linux was the operating system for a variety of networking and multimedia applications.”

Another leading embedded-platform vendor, Wind River Systems, made the move to an embedded-Linux technology-based platform for network equipment in early 2005, followed a few months later by the release of a Linux technology-based platform targeting mobile phones, set-top boxes, personal-video recorders and other consumer-electronics devices. “We hit critical mass with Linux just in the last six months, in terms of having a full range of products that meet a lot of different markets,” says Glenn Seiler, director of Wind River’s Linux product line. “We’ve been supporting Linux in the networking space for two years, and now we’re making a new push for Linux in the consumer space. We have some big customers, like Nortel, that are putting Linux into core and edge networking devices.”

Embedded vendors have been accelerating efforts to make Linux accessible and easy to use, says IDC analyst Fred Broussard. He says embedded-Linux vendors are “creating a developer kit that’s easy to use, virtually giving it away, then partnering like mad on the hardware side and on the software side so that developers have reasons to use the kit. Before, they would sell their kits for thousands of dollars.”

Toasters No More

To a number of industry observers, the growth of Linux in the embedded space means more than simply having a new OS in town. Linux’s arrival is a sign that embedded systems and devices have matured to enterprise-class standards, and, as a result, breaks down the wall that separated the devices from enterprise systems. The Linux that’s used within the devices is the same OS that runs the largest servers.

“Embedded devices used to be unconnected, with proprietary or simplistic user interfaces,” says Weinberg. Now, he says, “The modern embedded application is a network device that presents a Web-based management interface. It has core functionality defined in terms of its connectivity to the Internet. Printers, multifunction devices and phones have

gone from being standalone proprietary peripherals to standards-based network devices. We're using an enterprise operating system in what was a toaster environment. It's very liberating."

The rise of Linux means devices can begin to incorporate a range of functionality far beyond the limited capability of the small-footprint proprietary OSs of the past. With Linux, embedded developers "can make use of a wide range of publicly existing device drivers, design systems using the latest communication protocols, more easily build robust graphical interfaces, supplement existing platforms with technology leveraged from the enterprise Linux domain, and enjoy royalty-free production licensing," according to Chris Lanfear, director of VDC's Embedded Software Group.

70 Kb is as 70 Kb Does

A major criticism is that Linux requires more system resources than many of the small-footprint real-time operating systems (RTOSs) specifically designed to run within limited-memory devices. "Embedded Linux is no different from desktop or enterprise Linux," Weinberg points out. "Traditional RTOSs were animals that lived in under 60 to 70 kilobytes. But they were relatively trivial pieces of software performing a yeoman's task in traditional embedded systems. But in terms of the size and scope, they're very small. And my feeling has always been '70 kb is as 70 kb does.'"

As essential features are added to an RTOS to make an embedded device fully functional, the size often swells beyond an equivalent Linux deployment. "Once you take a traditional RTOS kernel, and you bolt on TCP/IP and a Java* Virtual Machine and a real file system, it's actually bigger than the equivalent Linux integration," says Wienberg. While Linux requires at least 1 to 2 MB, Weinberg adds that much of the device hardware becoming available these days has that capacity and more.

Another concern raised with Linux is that it's not inherently an RTOS—one of the leading features often associated with embedded devices. "In addition to its substantial code size, Linux isn't innately well equipped to manage system resources or handle strict real-time requirements as efficiently as smaller, more nimble embedded and real-time operating systems," notes Lanfear.

However, Weinberg questions the perception that most embedded systems need to be real-time. "Linux is not a real-time operating system, but many real-time operating systems are overrated," he said. "Maybe only 10 to 15 percent of embedded applications have a real-time requirement. The association between real time and embedded has always been overstated." The rise of Linux within embedded environments may underscore this fact as well, Weinberg continues.

Small-footprint RTOSs are unable to handle the demands of increasingly sophisticated and enterprise-aware devices, either. For example, many devices—particularly smart mobile phones—"have a software content that's doubling every year," says Weinberg. "They now have between seven and ten million lines of code. That much code overwhelms the classic RTOS. You end up with an inverted pyramid, where the application is totally out of scale with the underlying RTOS, which becomes incapable of managing the application's complexity. Like enterprise workloads, embedded applications are becoming huge, complicated systems with thousands of threads and great complexity."

Ripple Effects

The advantages of embedded Linux are most likely to be directly seen first by device manufacturers and developers. However, end-users and their organizations also will feel the effects of enhanced functionality, more competitive pricing, and faster, more responsive innovation as vendors cut loose from proprietary systems and migrate to embedded Linux.

Price is the advantage many people are aware of with Linux deployments, since developers and users can download and deploy as many copies of Linux as desired without additional licensing fees. Industry experts caution, however, that Linux isn't free, since new products require investments of development resources. However, the flexibility and streamlining that Linux brings to bear can ripple through the supply chain. Device manufacturers and end users "go from a royalty-bearing proprietary operating system, or in-house operating system with amortized support cost based on internal personnel, to a community-supported and royalty-free based operating system," according to Weinberg. "It's an easy transition in terms of improving the margins on bills of materials." This savings, of course, is translated to the retail costs of embedded devices.

"From the perspective of the appliance builder, using Linux or using open source reduces the cost of goods enormously," says Bernard Golden, chief executive officer for Navica and author of "Succeeding with Open Source." "They don't have to write it, or build it; they can just piggyback on it. Maybe they add some secret sauce on top of it, or maybe modify it somewhat for their purposes. It certainly reduces the costs, which means the product can be delivered at a much lower price point, which means that more people will adopt it."

Since the Linux source code can be acquired at no cost, the testing and prototyping phase of development is greatly accelerated as well, says Wind River's Seiler. "Linux gives developers very rapid prototyping, with access to source code at no cost. Therefore, developers and manufacturers have more control over their projects."



“Printers, multifunction devices and phones have gone from being standalone proprietary peripherals to standards-based network devices. We’re using an enterprise operating system in what was a toaster environment. It’s very liberating.”

—Bill Weinberg

Faster Time to Market

Fast time to market is the ultimate result of the control and flexibility enabled through embedded Linux. “Device manufacturers need fast time to market, and this requires more control and flexibility over their projects,” says Seiler. Since many communities and projects are occurring around Linux, new innovations are constantly at the ready. Time to market is no longer dependent a new release or update from the operating system vendor. “There is more innovative technology being developed in open source than any one company, even Microsoft, can deliver by itself,” he adds.

Such ongoing innovation can make a difference in a global marketplace. This has been critical for Motorola, which has been incorporating Linux within many of its mobile-phone products for the past three years. “The selling seasons for our cell phones peak at different times across the world,” says VandenBrink. “The Chinese New Year is a big time for selling phones in China,” he explains. “Christmas and back to school are big times for selling in the United States. Depending on which region and country you are, you have to be able to release devices with needed feature sets at varied times throughout the year. Linux greatly enhanced our ability to make changes when we needed to make changes to address certain parts of the global market.”

With proprietary vendors, VandenBrink said, new releases are limited because each company has its own release schedule. “With Linux, we don’t miss a turn in the market because a particular piece of software was ready or wasn’t ready. It puts more of the control back into the handset manufacturers’ hands.”

In the process, manufacturers or vendors can concentrate on their core applications and functionality, without the distractions of OS design. “They don’t have to spend a year-and-a-half designing plumbing,” says Golden. “They can focus on what they can add on top of that, to that differentiate themselves.”

Wind River’s Seiler agrees “most companies don’t want to be in the operating system business. They want to be in the business of providing cell phones, routers or TVs.”

Enhanced Innovation

Low costs and outstanding flexibility helps increase the innovation rate possible in the embedded-device marketplace. “We went with Linux because it provides a rate of innovation we found compelling,” explains VandenBrink. “When we looked at Linux, we saw the ability to push things strategically faster along a certain path, without having to wait on, or be gated by, somebody else’s release schedules. Linux provides a way to rapidly innovate and get new features into a device with a lot more freedom.”

With Linux under the covers, device developers and manufacturers can modify functions as needed. “One company I worked with needed a near real-time operating system, so they were able to use a Linux that used a different memory mechanism to help ensure quicker response times,” says Golden. “Linux is also very useful to embedded or special-purpose hardware manufacturers, because they can modify the code to enable support for their specific hardware. With Windows, by contrast, manufacturers have no control over the software, so they have to wait for Microsoft to put support for their hardware into the operating system. With Linux, manufacturers can modify the operating system and ship that version with their product.”

Weinberg observes that Linux has evolved embedded systems to become “peer devices—on a par with the desktop and a peer of the server,” he explains. “They can no longer afford to rely on simplistic ‘bare-metal’ RTOS models. There’s no longer a discontinuity between embedded and enterprise. They demand an enterprise-class platform underneath. Linux, whether it’s on an enterprise server or on an embedded device, is a single OS with a continuum of applications.”

Worldwide Shipments of Embedded Linux* Operating Systems

2003	\$73.2 million
2004	\$91.5 million
2005	\$133.6 million

Source: Venture Development Corporation

Joe McKendrick is a research consultant specializing in IT and organizational development trends, authoring special reports for IDC and Gartner. He’s contributing editor to Database Trends & Applications and editor of Data Center Trends & Applications. Joe can be reached at joe@mckendrickresearch.com.



The Ticket to IT Efficiency

With its Linux foundation and support from IBM and SAP, Ticketmaster transforms its financial-systems environment **BY NEIL TARDY**

As a growing company, Linux* made sense for Ticketmaster (www.ticketmaster.com). Now, as a global operation serving more than 9,000 clients worldwide, Ticketmaster continues to rely on the cost-effectiveness and flexibility of Linux.

In short, Ticketmaster is finding that Linux makes more sense than ever. The company's in the final stages of an initiative to revamp its financial-systems operations. With the help of IBM and SAP, Ticketmaster is replacing three regional datacenters with two new sites—a primary site located on the West Coast and a disaster-recovery (D/R) site on the East Coast, both using IBM* System x* hardware.

“The ticketing marketplace is continuing to evolve, and Linux and the SAP platform provide Ticketmaster with the flexibility to bring new services online quickly,” notes IBM's Steven L. Canepa, vice president, global media and entertainment industry. “The Ticketmaster solution is an excellent example of the role Linux is now playing in transforming core business-solution environments.”



Like many companies that conduct a significant percentage of their business online, Ticketmaster's ongoing growth—including the dynamic products and services offered to its clients to help the company sell more tickets and more effectively manage its ticketing operations—presents continued challenges for its IT department. Ticketmaster had already conducted various IT consolidation projects in recent years,

like general ledger and accounts payable/accounts receivable.

“Initially, we were a smaller company, and we had much more of a regional focus,” explains Ticketmaster's Vivek Kalra, vice president, financial systems. “But as we were expanding and growing across all areas of our business, it was just a natural evolution to look for efficiencies. We were taking the next step.”

Transforming the Silos

The next step, as Ticketmaster saw it, was to further strengthen its relationships with its clients—the event organizers and promoters, many of whom also run the stadiums, arenas, concert halls

and other myriad venues that contract with Ticketmaster to handle ticket sales for their events. Given the company's mission to help its clients sell more tickets, creating efficiencies was paramount. By updating the financial systems,



But as we were expanding and growing across all areas of our business, it was just a natural evolution to look for efficiencies. We were taking the next step.”

—Vivek Kalra, vice president, financial systems, Ticketmaster

Taking the Next Step

Ticketmaster is an almost ubiquitous presence in today's consumer world. The company, which has its headquarters in West Hollywood, Calif., provides exclusive ticketing services for hundreds of leading arenas, stadiums, performing-arts venues, museums and theaters. Conducting sales online, over the phone and at more than 6,500 retail Ticket Center outlets, Ticketmaster sold 119 million tickets—valued at \$6 billion—in 2005.

Founded in 1976, Ticketmaster has undergone many changes and transformations in its history. But one constant seems to be growth. From 2000 to 2005, Ticketmaster increased the number of tickets sold on the system from 83 million to 119 million. In that same timeframe, revenue rose from \$519 million to \$950 million.

when in 2003 it turned its attention to the area of financials. Looking first at its operations in the United States, the company felt it had outgrown its current setup. In the United States, Ticketmaster was running three datacenters, all housed with different systems that ran different basic accounting applications

Ticketmaster felt it could better manage the critical process of settling the funds it receives from ticket sales and directing the appropriate contractual percentages to its huge roster of clients.

“We process an incredible volume of ticket sales on behalf of our clients,” says Kalra. “Every financial transaction



must be settled in an accurate and timely manner. It's just mission-critical to our business, so we really needed a stable platform that we could rely on."

With the three regional datacenters serving the eastern, central and western portions of the United States, Ticketmaster essentially had three teams handling this process for three different sets of clients. As noted, each datacenter had its own IT infrastructure and application environment. The need for greater uniformity among the datacenters was becoming apparent to Ticketmaster's internal staff of users responsible for calculating client settlements.



Recognizing that it needed to provide the same user experience to everyone working with its financial applications, Ticketmaster began investigating ERP solutions. By using a packaged

"Our existing applications needed to be updated to talk to each other in order to streamline support issues, training issues—we needed a system that would bring it all together in an integrated environment instead of within silos," says Kalra. "When someone required a consolidated report, data would need to be extracted and run in a centralized warehouse, or sometimes manually in Excel."

application, the company was breaking new ground. Ticketmaster has traditionally relied on—and continues to prosper with—home-grown applications. That the company has accomplished so much with its in-house skills is to a significant degree a result of its long-time use of Linux. Ticketmaster has deftly leveraged the platform over the past five years, converting many applications from other operating systems (OSs) to



SAP and its mySAP Business Suite. SAP's support of Linux is well-established—the Waldorf, Germany-based company, in fact, was the first to run mission-critical ERP applications on Linux.

While Ticketmaster utilizes several of mySAP's financial modules, it was drawn to one particular module, called Incentive and Commissions Management (ICM). ICM allows Ticketmaster's

will keep pace and support every impact to the financial systems and processes.

The staged rollout began in spring 2005, when the company's West region began using the general financial modules. By November, Ticketmaster's three U.S. regions were all up on the general modules, while ICM was live for the West region. In June 2006, the company completed the ICM implementation in two of the three U.S. regions. Ultimately, SAP will be in place worldwide at Ticketmaster—a rollout is already underway for the company's Canadian financial-systems group.



We've found that other big companies are now running Linux as the main platform. Linux has always allowed us to run a cost-effective hardware infrastructure."

—Jorge Chereque, business systems manager, Ticketmaster



Linux systems host the mySAP financial modules in use at Ticketmaster. "When we went through the vendor-selection process, we were focused on the functionality, and that's where SAP came up as the frontrunner," says

Linux. This is directly related to Linux's maturity, the greater amount and variety of software available, and the depth of experience within the Ticketmaster organization. Currently, with the exception of databases, the company's core back-end applications run on the Linux platform.

"Linux is leading-edge technology," Ticketmaster's Jorge Chereque, business systems manager, points out, "and we have very good Linux expertise in-house."

Adds Kalra: "We're pleased with our commitment to Linux. We believe in the value."

An Efficient Environment

In an ERP solution, Ticketmaster was seeking enhanced functionality for its end users. But, given the company's commitment, the capability to support Linux was equally essential. Thus, in late 2003, Ticketmaster settled on

financial systems team to more-easily manage client accounts, providing an integrated solution that's flexible enough to meet the company's increasingly complex calculation requirements. As the company adds new ticketing products and components in response to continued growth and to better serve its clients, it needs an IT environment that


Kalra. "But if Linux was not certified with SAP, that certainly would have been a large factor in the decision-making process."

Just as Ticketmaster's IT group knows Linux, it knows IBM. Chereque notes that thousands of System x servers are deployed throughout the company. However, at the three U.S. regional centers, the financial systems were

Microsoft* platforms. With the move to the primary and D/R sites, Ticketmaster replaced its Windows* systems with new IBM hardware running Linux. The company has dozens of servers assigned to SAP, but more importantly, it now offers a single platform that helps its business operate more efficiently.

"From a hardware point of view, it was easy to think of IBM, because of the relationship we've established with them," Chereque adds.

IBM not only provided the new systems at Ticketmaster, it lent its



considerable consulting expertise during the SAP rollout. IBM is an acknowledged leader in supporting business-critical applications like SAP's. From their global-alliance relationship that spans nearly three decades, SAP and IBM share more than 10,000 customers worldwide.

"Without prior experience with implementing ERP, we were looking for a provider that could mitigate some of the risks of going with an ERP package," Kalra explains. "Also with the fact that ICM is a relatively new module, we needed a strong implementation provider that had a healthy relationship with SAP. We relied on some of the specialists from both IBM and SAP when it came to fine-tuning and configuring."



CUSTOMER: Ticketmaster

HEADQUARTERS: West Hollywood, Calif.

BUSINESS: Global ticket sales

HARDWARE: System x hardware

SOFTWARE: SAP's mySAP Business Suite

CHALLENGE: Improving financial systems management

SOLUTION: Using System x hardware and SAP software to consolidate and improve availability of IT environment

nearly complete in the United States, Ticketmaster is better than ever. In addition to delivering a better environment for its end users, the project is helping the company further strengthen the relationships it has with its clients. Ticketmaster clearly recognizes the value of Linux to its operations.

"We've found that other big companies are now running Linux as the main platform," says Chereque. "Linux has always allowed us to run a cost-effective hardware infrastructure."

Neil Tardy, a freelance writer and editor, has covered technology since 1996. Neil can be reached at ntardy@msptechmedia.com.

Linux is the Foundation

With its transformation of the financial systems team's IT environment

The Accidental Advantage

Implementing Web services with Linux on a System i platform provided Olympus America with unexpected benefits **BY KRISTIN LEWOTSKY**



Andrew Bartellino, manager, technical support and operations, iSeries hardware, Olympus America Inc., chose Linux on System i hardware to increase capacity and improve system performance.

OLYMPUS

Serendipity can be a wonderful thing, especially when it results in saving time and money while achieving improved performance. When Olympus America needed to scale its Web-services effort, the company activated an IBM* System i* partition fitted with Linux*. Going with a Linux implementation was the fastest way to get the capacity it needed. The group didn't go looking for high performance and reliability,

which made the 3:1 performance enhancement the Linux solution provides even sweeter.

If you think of Olympus as a manufacturer of cameras and other consumer electronics, you'd be partly right; that's far from the whole story. Olympus America and parent company Olympus Corp. have leveraged their imaging expertise to penetrate the healthcare market with diagnostic, therapeutic and imaging equipment, such as endoscopes, blood analyzers and confocal microscopes. In fact, healthcare makes up the majority of the company's business.

When you play in the healthcare market, lives can hang in the balance. Providing and servicing equipment as rapidly as possible is key. After implementing JD Edwards World enterprise resource planning (ERP) software on the i5/OS* platform, the company turned to a Datasweep solution to service and repair medical products. While it achieved its, the solution had one significant downfall—it had no mechanism to access the core data it required from the ERP system.

Some of the access was periodic, but some of it needed to take place in real time. The service and repair system for the medical division, for example, needed real-time pricing information as well as service-contract information.

To solve the problem, the team turned to a Web-services approach (see the sidebar, "Understanding the Jargon," page 25). The Olympus America IT group built more than 50 different Web services, leveraging Java* 2, Platform Enterprise Edition (J2EE) and Apache

Tomcat server, including price lookup, inventory allocation, bill-of-material retrieval and service-contract retrieval—as well as a credit-card authorization that wasn't a function of the ERP system. Initially, they implemented the applications on a Windows* server. Serendipity came into play when they needed additional availability.

Olympus purchased a System i platform with 32 processors, 24 of which were active. "IBM had an offering that gave you the opportunity to use Linux as part of your system," says Andrew Bartellino, manager, technical support and operations, iSeries* hardware, at Olympus America Inc. "When we purchased our 890, it was purchased as a Standard edition. We then upgraded it to an Enterprise edition, and with that there was the promotion of a free processor to be set up in Linux." The team took advantage of that offer, creating a partition using one of their current processors and installing Linux. "It really was an offering from IBM that gave us the opportunity to look into [Linux] and further explore how it would integrate and work within our organization," Bartellino adds.

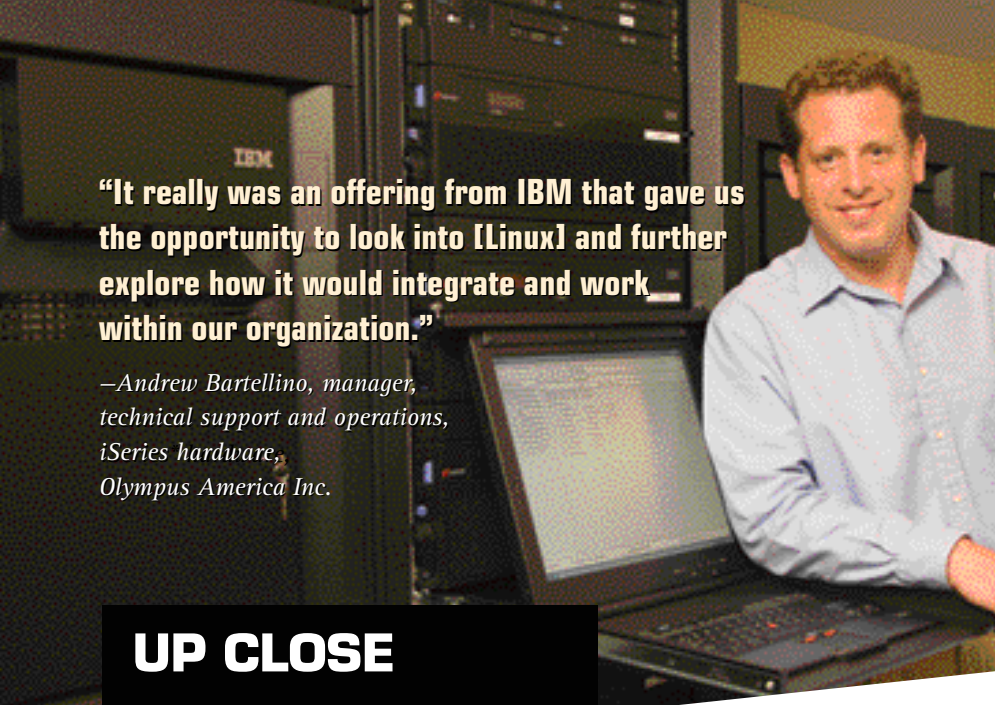
"Basically, it was a question of availability," says Director of Systems Development Bob Heiler. "We knew we had the partition out there, so we said, 'Let's take advantage of this, as we don't have to go out and buy another server. Let's give this a shot.'"

New Territory

It's human nature to want to stick with the tried-and-true solution, especially



PHOTOGRAPHY BY JIM ABBOTT



"It really was an offering from IBM that gave us the opportunity to look into [Linux] and further explore how it would integrate and work within our organization."

—Andrew Bartellino, manager, technical support and operations, iSeries hardware, Olympus America Inc.

UP CLOSE

CUSTOMER: Olympus America Inc.

HEADQUARTERS: Center Valley, Pa.

BUSINESS: Consumer electronics

HARDWARE: Two IBM System i5 model 570s

SOFTWARE: Lotus Domino, Lotus Notes, Red Hat Linux, Tivoli Storage Manager, JD Edwards ERP software, Java 2, Platform Enterprise Edition, Apache Tomcat server, Java Development Kit

CHALLENGE: Consolidating its IT environment

SOLUTION: Running a Linux partition on System i hardware, which doubled the workload they could previously perform on a Windows server

when access to company-critical systems is at stake. As a result, organizations that haven't run Linux are often uneasy about making the leap. Yet, as Olympus America's experience shows, the shift shouldn't be a big, scary thing. For Paul Franceschini, manager of systems development, and Narayanan Pillai, project leader, the chance to work with Linux was a long-awaited professional opportunity.

"Pillai and I knew that there was a lot of success with the Linux OS (operating system) running on a System i partition," Franceschini says. "One of our career goals was to do this someday. As

soon as we were presented with the opportunity, we worked it out."

The Olympus team had an advantage in Pillai, who'd spent several years exploring Linux on his own time. "Whenever I've played in the J2EE arena, it's always been on a Linux server at home—and one of the outstanding things that I've seen in that is that both Apache's Web server and their application server Tomcat run beautifully on the box," he says. "We knew that Linux would run extremely well on the System i hardware. When Andrew said they had a guest partition available, we jumped at it."

To get the partition operating, the team first had to install the OS. They chose Red Hat Linux. Next, they installed IBM's 64-bit Java Development Kit and then installed Apache's HTTP server and Apache's Tomcat Application server. "Once we had Linux set up and running," says Pillai, "Apache's binaries for Linux Web server and Tomcat were so simple to set up that we could do it blindfolded, almost."

Because of the nature of Web services, the change of OSs from Windows to Linux was essentially transparent to the Web-services framework and applications the group had built. Overall, the implementation process went smoothly.

"I think the bottom line is that although we did have some good in-house knowledge, we feel that someone who has a basic knowledge could be successful with this same sort of environment," says Franceschini.

More Bang for the Buck

The IT department's decision to go with Linux was purely pragmatic, but it yielded unanticipated benefits. The JD Edwards ERP solution was resident on a System i model 890 with the Linux partition. The group had already been running the Web-services framework on a separate non-IBM server using two processors and 2 GB of memory. Once they activated the Linux partition on the System i platform, it performed the same work with only one processor and 1 GB of memory.

When Pillai and Franceschini began experimenting with the capabilities of the Linux partition, they made a startling discovery—they could double the load on the Linux partition without any change in performance. "It was serendipity more than anything else," says Pillai. "Once we had the environment up and running, just for giggles we sent twice as much work to the Linux partition as would go to the Windows server. The response times didn't change." At maximum, they were sending three requests to the Linux box for every one request sent to the Windows box, for a 3:1 performance improvement. They've since scaled back from that point.

There are mitigating factors, of course. The Linux partition has the advantage of residing on the same system as the ERP application, as opposed to a separate server. The Linux partition thus interfaces with the data over a Gb internal virtual LAN, compared to the Windows server, which must access data over a physical LAN operating at 100 Mbps. "None of this should have made as much of a difference as we got out of it," Pillai says. "The [System i] hardware was definitely a factor. Definitely, Linux is also a factor, and of

course the [Java Development Kit] could be a factor," he adds.

The process wouldn't be complete without a few challenges. The limited availability of Linux training is one, Franceschini says, though there are more offerings daily.

A somewhat bigger problem arose from the fact that the group is running both Linux and Windows versions of the Web-services environment and vendor applications. Part of system administration is monitoring processes and performance. That's typically straightforward with the right tools, provided you're working with a single OS. Working with two OSs can get tricky. "This Linux environment is load balanced with a Windows server that runs the exact same applications," says Franceschini. "The challenge there is to find tools that allow us to monitor both of these operating systems at the same time." Currently, the team manually monitors status; they're still seeking a solution that can do both.

The dual-OS environment poses other challenges. "Something as simple as deploying the application to both servers requires two different processes," says Pillai. "For Windows, we just map a drive and drop it, and that's the end of it. With Linux, because we don't run Samba, we have to actually have a client on the deploying machine to do a secure FTP transfer to the Linux box."

Going Forward

The group chalks up a significant portion of their Linux success to the System i platform. "It's important to understand that we did this on the System i platform," says Heiler. "We got a lot of support from IBM trying to run Linux, and it was successful because it just fit very well."

"Certainly, the reliability that the system provides is great," Bartellino says. "Our downtime is very limited on this platform unless there's some critical fix we have to perform." Like Heiler, he considers the Linux/System i synergy a crucial part of the value proposition. "I think the marriage between Linux and the System i platform is really, really good. Running it on a different platform, would it perform as well, would it be as reliable? I'm not 100-percent sure," he confesses. "We can certainly speak on behalf of the System i solution itself—its superior hardware and the reliability is unbelievable."

pretty much took the Linux partition, saved it and restored it using IBM Tivoli* Storage Manager (TSM)," says Bartellino. "The integration for that piece was seamless. It took maybe a few hours at best."

In addition to the obvious performance advantages, the Linux effort has given Olympus the flexibility to pick and choose OSs going forward, depending on the application. If the partition and the data reside on the same System i platform, for example, the decision is straightforward. The company has no plans to migrate to either all Linux or all Windows, but will continue to run a mix of the two.

Understanding the Jargon

Web services: Web services refers to a platform- and language-independent method of passing data and sharing subroutines via the Internet. Using Web services, an IT team can write a single data-retrieval routine, for example, that will be used by applications from multiple vendors, written in different languages and running on different operating systems (OSs).

Partitions: To optimize performance, servers can be partitioned. Depending on hardware and OS, servers can be split into physical partitions or logical partitions.

Apache Tomcat: Apache is open-source software that allows ordinary servers to act as Web servers, loading Web pages, etc. Apache Tomcat works with Java* servlets that enhance Web-server functionality and allow the server to interface with other systems.

J2EE: Java 2 Platform, Enterprise Edition (J2EE) is a software tool for building Java applications, for example for services. IBM provides several different types of Java Development Kits (JDKs) to simplify the process.

JDKs: JDKs are designed to simplify creating applications in the Java environment.

—K.L.

The model 570s run a variety of workloads, most prominently i5/OS with five partitions running JDE, Lotus* Domino*, Lotus Notes* and SPSS. The IT department also has four native AIX* 5L partitions running TSM and one Linux partition. Linux isn't currently a factor in their high-availability (HA) solution. They have, however, discussed creating a second logical partition on the second i570 that would initially be used for testing but could be used for HA in the future.

As to Linux, the team is pleased with their results so far. "It definitely worked out very well for us," says Franceschini. "The advice I would give to a company

that wants to [try Linux] but may be afraid to do it is to take an application that's not mission critical and just not be afraid to give this a try. It is easier than you first think to implement, and the reward is great."

Kristin Lewotsky is a freelance technology writer based in Amherst, N.H. Kristin can be reached at klewotsky@chezhardy.com.



ILLUSTRATION BY PHIL MARDEN

Scale, Stretch, Flex and Repeat

Nationwide's System z and Linux team create a flexible IT environment

BY AARON DALTON

By any measure, Nationwide is a successful company. The Columbus, Ohio-based insurance and financial services organization is among the largest of its kind. A member of the Fortune 100, Nationwide has more than \$158 billion in assets and nearly \$20 billion in revenue.

Nationwide's growth and success have brought demands for a high level of IT support. Customer data must be managed, Web sites and payment systems maintained, and communications secure and reliable. To keep Nationwide computing smoothly, the firm employs thousands of people in its IT department. "In our business, information is key," says Guru Vasudeva, AVP and Enterprise Chief Architect at Nationwide. "It is the core of everything we do."

Talented and reliable employees form one element of a strong IT department, but computer hardware is also important. As

Nationwide continued to grow and its computer applications proliferated, its IT environment became more complex. The company ended up with thousands of servers, many of them underutilized.

They weren't necessarily sitting idle, but they were called upon to use only a fraction of their computing capacity. An internal study found that 78 percent of its servers had an average utilization of only 10 percent.

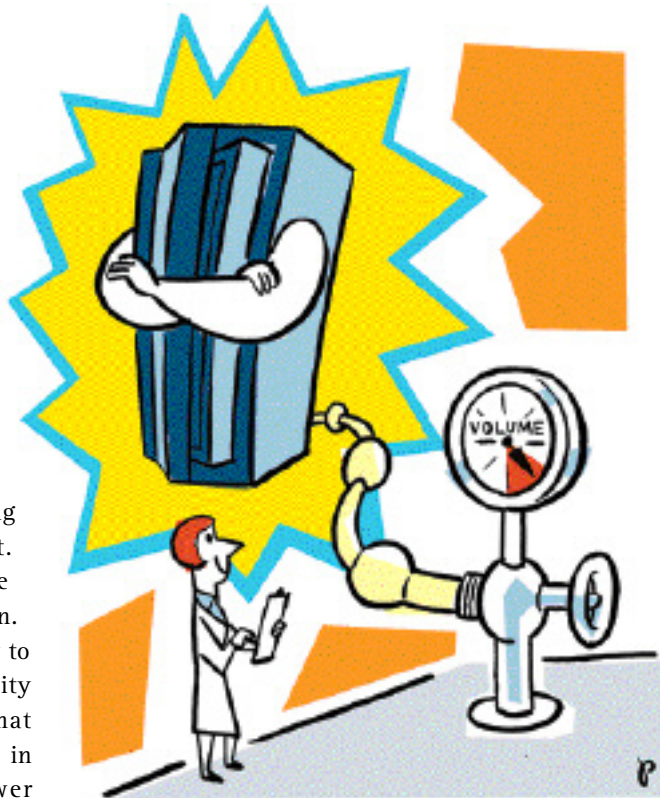
Vasudeva compares the situation to that of a 100-lane highway where average traffic occupied only 10 lanes. An individual server might see occasional utilization spikes, but often a vast amount of computing power sat idle. This idle capacity had its costs. IT personnel were required to set up, maintain and replace underutilized servers. The servers not only cost money to purchase and maintain, but they also occupied

costly floor space in an expensive datacenter. "Cooling, electricity and maintenance costs are quite high in a datacenter," explains Vasudeva. "If we had not controlled growth in the number of servers we used, we would have had to significantly expand our datacenter—and we did not want to do that."

Fortunately, Vasudeva's team had an alternative to expanding the IT department's footprint. Instead, they guided Nationwide to a virtual computing solution. "What we wanted was the ability to share all our idle server capacity across various applications that have different peaks and valleys in terms of computing power demands," says Vasudeva. "A technology called virtualization, combined with the Linux* operating system (OS), allowed us to achieve our goal."

Although virtualization sounds like a fancy concept, Vasudeva says it's really the revival of an idea that has existed for 40 years as proven technology in the form of mainframe computing. Nationwide purchased two IBM* System z* model 990 servers that run Linux. The System z platform allows Nationwide to place hundreds of virtual servers onto a single mainframe computer. Each of the virtual servers dynamically receives an appropriate amount of computing capacity as needed.

The move toward server virtualization hasn't only increased efficiency from a usage point of view, it's also increased average utilization of virtual servers to more than 80 percent. It has also brought significant financial benefits to Nationwide. Costs of Web hosting have plummeted, datacenter floor space has been saved and power demands have decreased to such an extent that Vasudeva estimates the move to server virtualization will result in \$16 million in savings over a three-year period.



Planning for the Super Bowl

Server virtualization has given Nationwide greater flexibility in planning for situations that call for brief spikes in computing power. For example, if Nationwide runs a television advertisement during the Super Bowl directing viewers to its Web site, the company knows millions of people will rush to its Web site (www.nationwide.com), requiring additional computing capacity to handle the demand.

In the previous server environment, Vasudeva's IT department needed months to purchase, install, configure, test and coordinate server hardware in preparation for one night's surge in Web-site traffic. With the on demand computing capability of the System z servers, Nationwide can essentially turn on new capacity as needed and turn it off when demand subsides.

Of course, the on demand capacity in a System z server isn't limitless. Vasudeva says that specifics vary depending on how a deal is negotiated, but that generally IBM ships the System z servers with spare capacity that can be activated for short time periods.

For example, a System z server may ship with 100 units of computing capacity, but an IT department may only need to use 50 of these units for its day-to-day operation. The System z computer knows how to divvy up those 50 units of computing power to the firm's 100 applications through a network of virtual servers.

Now the corporation plans to run a television advertisement during the Super Bowl. The IT department estimates the response to the advertisement during the show and in the week there-

after will require another 20 units of computing power. The company simply pays for a week's worth of 20 additional units of computing power up front, and IBM increases the System z machine's computing power for the agreed amount of time.

Vasudeva compares the flexibility of the System z platform to a cell-phone plan. "If you paid exclusively by the minute, you would have too much uncertainty," he says. "So you buy a 600-minute plan and get some predictability. Then, if you run over the plan, you just pay for those additional minutes." In contrast, with the traditional server environment, once you ran up against your minute limit, you would need to go out and purchase a new phone if you wanted to continue talking.

Even if computing demand increases beyond initial expectations, a System z user may not need to add new servers, but may be able to add components (e.g., CPU, memory) into the same System z machine. "You just cannot get this kind of sophisticated on-demand computing anywhere else," says Vasudeva.

The Linux Advantage

So far, Nationwide has placed the company Web site, some portions of its agency-management system, the employee-absence-management system and some smaller departmental systems into the Linux/System z virtual-server environment. The company plans to put significantly more of its Internet functionality onto this platform.

The move to the System z platform represents a major increase in Linux use for Nationwide. Previously, the company had only employed Linux in a limited way over a three-year period, but Vasudeva says his earlier experience with Linux had made him comfortable using the OS as an enterprise platform.

Moreover, Vasudeva points out that Linux offers its users certain advantages. For one thing, he says it's becoming easier to hire skilled Linux programmers. Vasudeva also praises the way in which Linux's open-source development model stimulates frequent advances in the OS's functionality.

From an IT department perspective, hardware neutrality is another major attraction. As it becomes commonplace for applications to run in a Linux environment, companies like Nationwide are no longer locked into certain hardware vendors and can choose whichever systems work best for their needs.

What Goes Around ...

Vasudeva sees server virtualization as a new-and-improved take on the tried-and-true. In understanding the reasons behind server virtualization, he believes that it's helpful to consider the history of how firms have tended to organize their IT departments over the past few decades.

Forty years ago, mainframes dominated corporate IT, but Vasudeva says that the older mainframes weren't flexible enough and were too tightly controlled by the datacenters. As computers became more affordable, companies moved to increase flexibility by giving each department its own servers. This decentralization seemed like a good idea at the time, until one day companies realized that they had 5,000 servers—each with its own maintenance and upgrade requirements.

To achieve enterprise quality, companies needed to move the servers back to the datacenter for effective management. That's when datacenter costs started to rise as companies struggled to cope with the maintenance, replacement and floor-space demands of so many machines. "Decentralization was cheaper on a unit-by-unit basis, but the total cost of ownership was more expensive," says Vasudeva. "We were making penny-


wise decisions that were dollar-foolish at the aggregate level."

In response, of course, companies like Nationwide have turned to server virtualization. Vasudeva points out that the concept of virtualization actually dates back to the 1970s, when computers were so expensive that departments and individuals had to share every last byte of computing power.

The return to the time-sharing concept of computing is taking place for a different reason now, he points out. "Where originally it was costs that drove us to use timesharing, now it is underutilization and complexity that drives us towards virtualization," says Vasudeva. The determining factor is no longer the price of computers, but rather the total cost of ownership (TCO) and complexity of managing and refreshing thousands of underutilized servers versus a smaller number of mainframe-type computers.

While individual servers must be refreshed every three to four years, mainframes don't need to be replaced as often. The solution that makes most sense is the one that Nationwide has adopted: to give each department or application the illusion of having its own servers, while really having them all share one computer that's smart enough to share resources based on demand.

The move to server virtualization means that the datacenter has been able to postpone an expansion even as Nationwide continues to grow. While the business has been growing and the IT department has been building new systems to accommodate and facilitate the growth, the bulk of that growth has been captured in the virtualized environment. Vasudeva says Nationwide has been able to defer at least \$20 million in expenses that it would have incurred



"Where originally it was costs that drove us to use timesharing, now it is underutilization and complexity that drives us towards virtualization."

—Guru Vasudeva, AVP,
Enterprise Chief Architect,
Nationwide

in datacenter expansion by now if it hadn't adopted a virtual-server model.

Efficiency and Flexibility

As older individual servers reach the end of their refresh cycle and go off warranty, he anticipates hosting the systems currently running on those machines inside the System z virtual environment. When that happens, not only will he be able to free considerable floor space—resulting in real-estate savings—he'll also reduce the time the IT department must spend in terms of server maintenance, cabling and replacement.

As the system becomes more efficient, it also becomes more flexible.

UP CLOSE

CUSTOMER: Nationwide

HEADQUARTERS: Columbus, Ohio

BUSINESS: Insurance and financial services

HARDWARE: Two IBM System z model 990 servers

SOFTWARE: Linux

CHALLENGE: Virtualizing its IT environment

SOLUTION: Using IBM System z hardware to run Linux, placing hundreds of virtual servers onto a single mainframe computer

Setting up a new server involved a process that once took weeks from start to finish—ordering, receiving, unpacking, cleaning up, setting up, etc. Now, the process of provisioning a new

server can be accomplished in days. "If you realize that you are running out of capacity, you just increase capacity," says Vasudeva.

A more-efficient computing model, lower TCO, greater efficiency, time freed to work on other projects, floor space liberated for other uses—however you look at it, the System z/Linux technology-based virtual-server environment has been a great boon for the Nationwide IT department and the company as a whole.

Aaron Dalton is a writer who specializes in business and technology topics. Aaron can be reached at aaron@imaginationwins.com.



0414: How to use `strchr()` in C++

Here's how to use the `strchr()` function in C++: `strchr("abc", 'c')` returns 'c' and `strchr("abc", 'd')` returns 0.

Let's look at a few examples:

```
#include <string.h>
using namespace std;
int main()
{
  char str[] = "abc";
  char ch = 'c';
  char* p = strchr(str, ch);
  cout << p; // outputs c
}
```

Then and Now

Tips on integrating open source as both an IT tool and a business enabler

BY JIM UTSLER

“**T**here’s a movement afoot. Soft murmurings of revolution and talk of fundamental change. A whisper of voices being whisked around the globe on the crest of an electronic wave. The clacking of keyboards laying bear to lines of code written in mysterious languages.”

That was how I began an article I wrote seven years ago. The topic? Apache, one of the most successful open-source projects ever (aside, of course, from Linux*). Things have changed quite a bit since then, and the “self-proclaimed hackers” I further mention in the article have largely given way to ranks of professional programmers often working for industry giants such as IBM.

And what was once considered a provocative movement—a way of sticking it to the man, if you will—has since morphed into an interesting new business model, with many IT businesses dedicating personnel to work collaboratively on open-source projects. Rather than endure the costly and lengthy process of building their own software solutions from scratch, they’re offering up their expertise to help innovate open-source solutions and then apply them to their own offerings.

As things have turned out, this has been a boon to customers, who now have more options than

ever when deciding which applications to introduce into their IT environments, whether open source or private source. Both Apache and Linux, for example, have since become ubiquitous, with IBM using Apache HTTP server as a core component of WebSphere* Application Server (WAS) and the company supporting Linux across its entire server and software portfolio.

But it’s because of the success of open source that users have to be more conscious about its uses and how to properly implement it. That’s why IBM has fielded a set of guidelines to help its customers with choosing, deploying and integrating open source as components of their everyday computing environment. The following is a high-level look at these open-source guidelines and how to apply them in practice.

Insist on Open Standards

“Open standards are the key to increasing choice and flexibility,” notes John Meegan, Software Group Open Source Strategy lead. IT environments based on open standards including those implemented with open source, private source, and mixture of the two, allow companies to choose solutions from multiple vendors and be assured that the different solutions will work together.

This is opposed to more proprietary solutions, which limit customer choice to a single vendor or a single tool, thereby limiting integration opportunities.

Open standards would include specifications for APIs and data and file formats that are openly documented, as well as protocols and development languages such as HTTP, HTML, TCP/IP, XML and SQL, which have evolved under the cooperative tutelage of organizations such as W3C, OASIS, OMA, ISO and IETF. Because these are open standards, they can span platforms and software solutions, making interoperability much more likely. "IBM has been a strong proponent of open standards for well over a decade," says Steve Gerdt,

rather because a company has a distinct business requirement for it. Questions to ask, according to Jollans, include, "How does it fit into your IT strategy? And how does that IT strategy help you achieve your business goals? You need to treat open source in the same way you would private source."

Some companies may have no need for open source at all. In fact, private-source operating systems, databases or application servers may be perfectly suitable for their purposes. As Jollans further points out, "There are some scenarios in most businesses where open source doesn't make any sense, and they shouldn't try to ram a

source/private source an either/or alternative." As he explains, companies need to look at both solutions and find the right fit for them. "They need to consider the pros and cons of each, taking into consideration functional robustness, performance, scalability along with service and support."

As part of this decision-making process, companies should combine their open- and private-source strategies. For example, if a company is looking at how best to build an e-commerce site, there should be processes in place that take into account both types of solutions, helping create a jointly built solution that uses the best of both open- and private-source solutions. "We follow that advice even within IBM," notes Jollans, "looking at where open source can play a role within our proposed architecture."

Another consideration is up-front costs versus recurring costs. Many people think that open source is essentially free, especially if it's downloadable and not part of a commercial offering. However, there are ongoing costs associated even with free open-source solutions, including employee training and perhaps the expense related to hiring new employees already familiar with the solution.

In the case of open source bundled with commercial offerings, support remains an issue. If something were to go wrong with a commercially available open-source solution, the company using it may have to have a contract with the vendor to provide one-off or even ongoing support. "A lot of people think open source is cheaper, and that's certainly true from the initial acquisition perspective," Meegan says. "But what customers need to understand is that the initial cost is a small percentage of the overall total cost of ownership. You have to take into consideration the cost impacts associated with migration, interoperability, ongoing services and support, and education and skills building."

"IBM has been a strong proponent of open standards for well over a decade."

—Steve Gerdt, Open Source Program Director

Open Source Program Director, "and we've seen that being locked into a proprietary business model simply isn't the best approach to business computing."

Similarly, open standards eliminate vendor lock-in. Companies can pick and choose from a number of vendors, knowing that if they move from one to another, their solution will remain essentially the same. "For example, in the case of Linux, you can go to Red Hat or Novell, or you can simply go to the Web and download it yourself," says Adam Jollans, Open Source Marketing Strategy Manager. "You're no longer tied into a particular vendor to support what is essentially your data and your IT infrastructure."

Evaluate Open-Source and Private-Source Options

As with any software acquisition, time should be taken to justify the use of open source. It shouldn't be adopted just because it's open source and "free," but

square peg into a round hole simply because they feel they have to keep up with the so-called latest technology."

However, other companies may indeed have a need for open source. In those cases, they have the option of downloading the software and trying it out before putting it into production. Although proprietary software vendors often offer trial versions of their solutions, open source often allows users to work on small projects, such as developing a simple Web-based application, without having to make that jump from trial version to production version and paying the associated costs. "You can easily download it from the Web and have your developers try it out without making a potentially costly decision on it," Jollans remarks.

Of course, going only with open source could itself cause issues, and, according to Meegan, "Very, few companies would consider open

Evaluate Open-Source Community and Maturity

"There are tens of thousands of projects that are open source," remarks Gerdt, "but I would hazard to guess that more than 95 percent of those wouldn't even be considered by an enterprise." This isn't necessarily because the open-source project hasn't been successful, but more because of the makeup or maturity of the community surrounding the project.

For one thing, companies want to be aware of who's involved in the project. If it's just a small group of well-meaning individuals, the solution may or may not be suitable for primetime deployment. It's much better to work with an open-source

robust and diversified and therefore more innovative.

"Just as you would with a commercial software vendor, you want to make sure the open-source community you're dealing with is mature and stable. You don't want it to go under only to lose support," Meegan says. "You want to look at the community to see how vibrant it is, and then look further down the line to see if there are multiple options in terms of vendors to provide the implementation services and support for that project."

As Meegan further points out, there are many fewer risks associated with

ing people willy-nilly implementing open source without proper authorization, an organization should educate its employees about exactly what open source is and how the licensing works.

This latter point is particularly important when people are considering acquiring open source, there are different types of licensing required, including dual licensing. In that case, there are licenses for end users and then licenses for companies wishing to use open source as part of a commercially available project. People should be made aware of the distinctions and the procedures required to comply with them. Aside from that, however, open-source licensing is relatively straight forward, especially when compared to private-source licensing.

"Some people think open-source licensing is complex, but it's really not," says Gerdt. "Of the 50 or so open-source licenses approved by OSI (Open Software Initiative), only 10 to 12 are commonly used, and once a company is familiar with them, there are little changes. On the private-source side, however, everyone has a different user licensing agreement and they can differ quite a bit."

Meegan notes that companies would be wise to have a legal review of all open-source licensing, despite the licensing issue being relatively straightforward. "There needs to be a formal review process in place whereby open-source licenses can be reviewed by a legal organization," he says. "In most cases, there aren't a whole lot of issues if the open source is being used internally. However, things may become problematic if a company is looking to include an open-source component into a product they expect to sell externally. In those cases, a thorough review of open-source licensing is absolutely critical." For its part, IBM has a steering committee that, according to Jollans, "does due diligence in terms of who owns the code, etc."

Additionally, private-source procurement procedures should also be applied

"Just as you would with a commercial software vendor, you want to make sure the open-source community you're dealing with is mature and stable."

—John Meegan, Software Group Open Source Strategy lead

community that is well established and well diversified. It should also have a transparent roadmap for future development, as well as a sustainable business model.

In many cases, the involvement of large ISVs and business partners can alleviate many concerns customers may have involving open-source projects. IBM, for example, is involved in a number of open-source projects, including Linux, Apache, Eclipse and Mozilla, lending not only developer support to the project, but also credibility.

Some open-source projects are essentially arms of for-profit companies, while others are truly open to all comers. The difference isn't only in motive, but also outcome, with the latter community being much more

working with an established open-source community. It's far less likely to dry up, leaving its open-source users withering in the sun without continued support. If a vendor does go out of business, there will typically be another to take its place, offering—thanks in large part to open standards—similar services.

Establish Open-Source Policies

Once a company has made a decision to go with open source, after having insisted on open standards, evaluated both its open-source and private-source options, and evaluated the communities with which it will be working, it needs to establish internal policies regarding the deployment and use of open source in the IT environment. Rather than hav-

to open-source acquisitions. After all, both are destined to become part of a business's mission-critical IT infrastructure, and thereby a key component of the business's larger strategic goals. These procedures should include the creation of a management-review system—if one isn't already in place—to help approve and track the use of open source both within the company and as part of an external effort to contribute back to the open-source community.

But as Meegan points out, some company employees see open source as a way to skirt internal procurement procedures. In fact, they may see free open-source software as a way to do an end-run around the entire process. "We've found quite a few IT departments that need a quick way to get started on a project, so they just download some open-source components," he says. "But then open source of all types start to creep

into the business, and all of a sudden, you're faced with interoperability issues because you now have applications running in the enterprise that the larger corporate IT department may not have known about."

By establishing and enforcing procurement procedures, and educating company IT teams, organizations can effectively avoid licensing problems and open source creep. Just as private-source solutions should go through management channels, so should open source.

Be Pragmatic

Open source is becoming increasingly popular, whether as direct downloads from the Web or as part of bundled packages. Whichever the case, companies should still work with the software as they would any private-source application, running proofs of concept and testing. Start small by picking a single

area where you think open source might be beneficial and plan on a staged implementation assuming the pilot results are satisfactory.

User should also make the decision to run open source based on both technical and business factors, with IT and strategic decision makers coming together to make sage decisions. After all, that software is going to become an important part of the larger tactical vision of the company. If open source doesn't fit, it shouldn't be forced. But if it does provide a business with benefit, then use it. In many cases, an open standards blend of open-source and private-source software will end up providing the best solution.

Jim Utsler, *MSP TechMedia senior writer, has been covering the technology beat for nearly a decade. Jim can be reached at jutsler@msptechmedia.com.*



Computing Like Life Depends on It

TGen teams IBM System p and System x servers with Linux technology to facilitate cancer-gene research

It seems everyone has known someone who computes like her life—or someone else's—depends on it. But at The Translational Genomics Research Institute (TGen), in Phoenix, where the study of billions of possible gene interactions can be compressed from 12 months down to one week, the impetus really is life. The scientists at TGen work to find patterns and interactions that can signal cancer-causing genes, which could lead to cancer-killing drugs. The faster they can comb through the haystacks of gene combinations for the desired patterns, the greater the hope for the eventual vaccine or treatment that will save a patient's life.

In research situations, it's important to have powerful computing ability on machines

that run well in an open-systems environment—sharing files and running parallel applications. Why open? Because researchers in this field are continuously writing new programs and upgrading old ones to quickly exploit their latest findings. A computer program to test a new theory that a type of cancer cell may express a particular gene or protein pattern could've been created on any number of machines. Rather than waste time and risk the possibility of conversion errors, open environments are the practical answer. That's why most researchers, who are usually fighting budget as well time constraints, write in C on an open-source system.

Open systems in use at the ASU-TGen High Performance

Computing Center, at Arizona State University (ASU), includes an IBM* System x* Beowulf Cluster consisting of 1,048 Intel* Xeon CPUs for parallel computations and several IBM System p* symmetric multiprocessing computer systems for memory-intensive computations. Typically, Beowulf Clusters are scalable-performance groups of usually identical PCs, running an open-source operating system (OS), such as Linux*. They're connected via a high-speed communication network and have common programming and inter-processor communication libraries that allow system resources to be shared for parallel computing.

"In the field of bioinformatics and computational biology, most open-source pro-

grams have been developed for use on Linux,” says Edward B. Suh, Sc.D., CIO. “Biologists are comfortable and familiar with Linux. Open-source programs are readily available; and open-source code can help optimize your computational requirements rapidly on your hardware.”

The center’s machines are linked using a network comprised of a Gb Ethernet connecting 524 nodes (each node has two Intel Xeon CPUs). Additionally, the center has 128 of these nodes connected with Myrinet, a low-latency, high-speed interconnect from Myricom. “This hardware provides the ability for the nodes to communicate

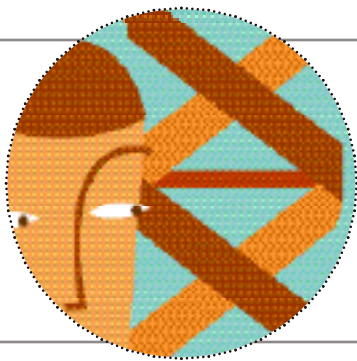
of Bill Gropp and Rusty Lusk; and Mississippi State University, under Tony Skjellum and Nathan Doss. IBM also made major contributions to development of MPICH, under Hubertus Frank.

Patterns, Associations and Interactions

To fully appreciate the impact of these technological contributions on life-saving science, you have to understand the processes. “We look at cancer,” says Suh. To be exact, the center’s genomic research encompasses tumor classification, risk assessment and prognosis, drug development, drug response, therapy development and disease progression. “We may be look-

bilized or attached to help them investigate genes’ relationships to cancer. By arranging many short sequences of nucleotides, A, T, C and G, such that they complement the actual genes’ sequences (for example, the complementary sequence to G-T-C-C-T-A will be C-A-G-G-A-T), scientists can measure the expressions of those genes over many thousands of DNA samples in a single experiment. They can then draw conclusions about the genes’ behavior (e.g., comparing diseased cells with healthy ones).

To isolate what genes are involved in a cancerous cell growth, scientists use a special type of DNA microarray chip, to detect whether a particular gene or set



“We are on the edge of an unprecedented quantum leap in medical science, dependent on the computational ability to perform complex data analyses and simulations as quickly as possible.”

—Edward B. Suh, CIO, TGen

much faster with each other,” says James Lowey, manager of the center. “This enables larger, more interdependent computational problems to be examined on this system.”

For example, a lab may want to perform a test that would require another lab’s research findings; perhaps statistically significant findings on diseased cells treated with a specific cancer drug, to see if that drug would be effective on another type of cancer with the same gene-expression pattern.

The ability to analyze these types of complex problems can be expedited by parallel computing with the use of MPICH, a popular message-passing interface (MPI) library, used in the Beowulf Cluster environment. MPICH was developed jointly by Argonne National Laboratory, under the direction

ing at melanoma to see what genes play a role or are important to this cancer. The patterns we deal with are not only based on how often something appears, but also the interaction it causes in other genes’ behaviors.”

“How often” is key. To establish those frequencies, Suh and Lowey provide the scientists with machines that can run analysis on billions of gene combinations.

“The sheer volume of data that is generated by these techniques requires a huge amount of processing power, and without access to this computational power it would take many years to analyze the amount of data being generated,” says Lowey.

The center, like other genetic-research labs, uses DNA microarrays, solid supports onto which the sequences from thousands of different genes are immo-

of genes is being expressed more or less than others, under given circumstances. This is called microarray-expression analysis. This technique can be used to determine the correct treatment for a disease. By comparing the similar expression of a gene pair in a particular form of cancer, it’s possible to compare the sample of diseased tissue from a patient and discover a match, thus making the diagnosis. This is why it’s so important to be able to compare billions of combinations quickly. Also, by isolating the expression pattern, expression chips can be used to develop new drugs.

To examine just 600 genes, in order to see the relationship of four gene combinations to other genes, scientists may have to look at more than 5 billion gene combinations. “If each gene combination analysis takes the range of 0.05

seconds, that's on the order of 10 years of computer analysis time," says Suh. "But if you have 100 CPUs working together, instead of 10 years, it would only take a month or two—or with 1,000 CPUs, less than a week."

"It's all about the compression of time," agrees Lowey. "If one problem takes four years to compute with one CPU, this cluster could enable it to run in just a couple of days." The computer runtimes will vary by application, of course, he adds. Some applications will have 90-percent parallel efficiency, while others may have only 50 percent. Parallel computing on a Beowulf Cluster enables the scientists to divide a compute-intensive large problem into small chunks that can be assigned to and run on many processors simultaneously. The level of parallel efficiency is based on the ratio of the additional increase in total computer runtime to the number of processors used.

The center is unique in that some of the microarrays it uses have up to 500,000 target spots on them. The data from the microarrays, which can build up very quickly, is stored in IBM General Parallel File Systems (GPFS). Application programs running on the Beowulf Cluster supercomputer draw from that data for analysis. To comb the data, researchers run it through a variety of computational algorithms, or chains of complex sets of orders.

"Most open-source scientific applications are developed in the Linux environment," says Suh. "Linux enables quicker development time."

One of the open-source applications that's frequently run on the center's Beowulf Cluster is the NAMD, a molecular modeling application developed by the University of Illinois at Urbana-Champaign.

The center may use the NAMD program to analyze proteins, which are important in cancer research because an abnormal protein signals that the gene that made it has mutated.

UP CLOSE

CUSTOMER: TGen

HEADQUARTERS: Phoenix

BUSINESS: Research institute

HARDWARE: System x Beowulf Cluster, System p hardware, General Parallel File Systems

SOFTWARE: Linux

CHALLENGE: Simplifying IT environment to help facilitate cancer research

SOLUTION: Using System p and System x hardware with Linux to tackle compute-intensive problems quickly

A protein is a three-dimensional structure encoded in the gene, which contains the DNA instructions on how that protein should be made up. Consequently, if this gene is mutated, the protein it produces will have an abnormal protein structure, produce abnormal hormones, etc.

To diagnose a protein, you need to unlock that structure. The NAMD program allows researchers to model the proteins by providing the tools to analyze the interaction between atoms, using the laws of thermodynamics. The researchers can observe how the protein moves by measuring the active and repulsive forces between atoms. Again, by comparing the behavior patterns of healthy proteins, the aberrant ones can be singled out.

If the protein structure is known, a drug can be developed that will attach itself to that shape. Without even touching the patient, researchers can computationally model the drug's interaction on a biomarker gene set—which is over expressed in cancer—to see if the new drug or an existing drug will reduce over expression and force the cancer into remission.

"What we would like to be able to do is to take a set of genes and say, this set has something to do with your cancer—to diagnose and prognosticate, and cre-

ate therapeutics for each cancer," says Suh. "For certain types of cancers, we are already doing that now. So we have some success stories."


But Suh isn't content to stop at these successes; he's asking other questions, such as whether the same cancer drug that works on breast cancer will work on other cancers. This requires complicated analysis. But, as the saying goes, "The impossible will just take a little longer."

Limitless Work

"The amount of work that can be done is limitless," he says. "We are on the edge of an unprecedented quantum leap in medical science, dependent on the computational ability to perform complex data analyses and simulations as quickly as possible. The science, the equipment is here. The key is processing and manipulating large volumes of data at extreme computational speeds." The ASU-TGen High Performance Computing Center appears to be on the right path for this race against disease, through the center's open network and operating environment; and its acquisition of increasingly powerful processors.

Maureen Nevin has covered information technology and finance since the 1980s. Maureen can be reached at MNevinDuffy@aol.com.

Chen Yu, vice president,
marketing, discusses why
YeePay chose to run Linux on
IBM OpenPower hardware.



With Linux on IBM OpenPower and BladeCenter JS20 servers,
China's YeePay starts small and grows

Turning Vision Into Reality

BY JIM UTSLER

Without visionaries spotting emerging business trends and markets, the global economy might not be what it is—or exist at all, for that matter. These imaginative individuals recognized, for example, that the Internet could create boundary-free ways of conducting business, allowing consumers to shop from their homes and businesses worldwide to reach them.

Of course, these same prophetic business leaders would be hampered if not for technology. It is, after all, IT that makes all of the above possible, with the Web allowing for customer purchase and payment (both private and commercial) and behind-the-scenes high-tech distribution systems providing same-day order processing. Without computers and associated technologies, even simple tasks such as balancing the books would still be a pencil-and-paper job.

One company that might be hailed as visionary, if for nothing other than spotting and addressing a weakness in an emerging economy, is YeePay, the Beijing-based provider of secure payment services for phone (both mobile and landline) and online users. Not only has it begun filling a gap in the Chinese economy, but also deployed a simple yet powerful IT infrastructure to make it all happen. Thanks to the use of an OpenPower* 720 server and two IBM* BladeCenter* JS20 servers, it's turning vision into reality.

Four Reasons

Based in Beijing and with offices in Shanghai, Guangzhou, Chengdu, Inner Mongolia and Silicon Valley, Calif., YeePay (pronounced "e-pay") is being positioned as

"the PayPal of China," according to Chen Yu, YeePay's vice president of marketing. But rather than focusing on credit cards, YeePay works with customers who use debit cards. This is in large part because the credit-card business in China isn't nearly as large as in other countries, especially those in the west. "Credit-card penetration in China is very low," Yu notes, adding that "The total number of credit cards is, I believe, around 30 million—and this is of a population of around 1.3 billion."

Cash- and bank- (or debit) card transactions, however, are much more prolific. Of course, cash can't be used for any online transactions (until, that is, someone comes up with a teleportation device), so that leaves bank cards to fill the void. But because bank cards require a PIN, there has to be a way to authenticate and process payments without users having to worry about PIN security issues. Hence YeePay, which works with both consumers and merchants to process transactions, whether over the Internet, via mobile phones or traditional landline phones. "China has become the largest mobile market in the world," Yu points out. "And it's the second-largest Internet economy in the world, with more than 100 million users. It also has nearly 400 million installed telephone units. There's an absolute need for our services, so people can securely and confidently make online transactions."

Since opening its doors for business in 2005, the company has already signed up more than 1,000 merchants to take advantage of YeePay services. Most of these are in the online gaming and online travel-agency industries, the latter of which Yu characterizes as the Chinese "equivalents of

PHOTOGRAPHY BY QUENTIN SHIH



“Now, with IBM, we have a total solution: the hardware, the database, the application server and even the technical consulting.”

—Chen Yu, vice president of marketing, YeePay

Travelocity and Expedia.” YeePay’s focus on these sectors is a prelude to a larger rollout, which is similar to the method by which credit cards became so ubiquitous in the United States. “Credit-card issuers started with restaurants and then hotels,” he says. “So right now, we’re working a lot with online-gaming companies and travel agencies, although we plan to branch out from there.”

Because the company plans to greatly expand its offerings over the next several years, it needed an IT infrastructure that could easily grow with it. This now includes the JS20 blades and the OpenPower 720, SUSE Linux*, IBM DB2* Universal Database*, IBM WebSphere* Application Server and payment-transaction applications YeePay built with Java* 2 Platform, Enterprise Edition (J2EE). Regarding Linux, Yu says, “We wanted to run a platform that’s based on industry open

standards. This gives us the flexibility to easily integrate other solutions and components as we need them.”

Initially, before beginning actual operations, YeePay had deployed a Sun Microsystems-based solution that included offerings from BEA Systems and Oracle. Although the company could have begun business operating within the confines of this environment, it chose to look at alternatives before going live. This, according to Yu, was because of four primary reasons—the first of which had to do with creating an IT infrastructure that came from a single vendor.

As Yu explains, “We didn’t want to have to deal with multiple ven-

dors. As it was, we had to deal with Sun for the hardware, Oracle for the database and then BEA for the Web application server. Now, with IBM, we have a total solution: the hardware, the database, the application server and even the technical consulting. It’s all integrated, coming from one vendor, and if something does go wrong, there’s no finger pointing; you know exactly where to go for support.”

The second reason had to do with “quality of services,” as Yu puts it, including availability, security and scalability. The latter was particularly important, because the company expects to grow well past its current 1,000-plus merchant customer base in short order. Now, using blades, the company can scale according to need, employing only the required capacity instead of too much, by, for example, purchasing a separate box that costs too

much and vastly exceeds current processing requirements.

“We wanted something that could grow with us, because we expect to see rapid growth of transaction volume,” Yu says. “So from IBM, we got business on-demand capabilities, and we have an infrastructure and platform that really scales.” He also notes that availability is crucial given YeePay’s business, which, he says, “requires 24-7 availability.” And thus far, the company has had no unplanned downtime.

YeePay’s third reason for moving to the IBM solution had to do with what Yu calls “good domain knowledge.” By this he means that IBM is already well entrenched in the Chinese banking industry, which makes it a good fit for the type of business YeePay is conducting.

As Yu elaborates, “Most of the major banks in China are IBM customers, and we think that’s a very important relationship we can leverage.”

The fourth, but no less important reason, simply has to do with customer relations. IBM has become a known brand in China, and according to Yu, “even average consumers recognize it, not just those involved in the IT industry. So our customers will see the IBM logo on our Web site and, by proxy, give us more credibility. That boosts our marketability.” And given that YeePay is one of the first companies to enter this industry in the Chinese market, this is a key advantage.

Although YeePay itself doesn’t have a datacenter (its IT assets are hosted by a third party), it does maintain and administer the systems. This allows YeePay to avoid the hassles of controlling the physical environment while still having control over the day-to-day operations of the systems. “We deploy the applications and we administer the servers ourselves. The datacenter host manages the network connections, the temperatures and other types of basic datacenter operations,” Yu says. This, he continues, “is a testament to the ease of use of the blades and the OpenPower system.”

Start Small and Grow

Without the forward-looking philosophy of companies such as YeePay, we might indeed still run in a cash-only economy, with trusted online payments not even being a notion. Instead, we now have the ability to make bill payments online, secure in the knowledge that our transactions will be well-tendered. But behind every great idea is often great technology, as the story of YeePay indicates.

By moving to an open standard such as Linux, YeePay can now continue to innovate, knowing that as its business grows, so can its IT environment. By simply plugging and playing, it has become part of the on-demand business world, avoiding the high overhead often associated with other solutions.

“As a start-up company, we wanted

UP CLOSE

CUSTOMER: YeePay

HEADQUARTERS: Beijing

BUSINESS: Secure payment-services provider

HARDWARE: OpenPower 720 server, two BladeCenter JS20 servers

SOFTWARE: SUSE Linux, IBM DB2 Universal Database, IBM WebSphere Application Server, Java 2 Platform, Enterprise Edition

CHALLENGE: Consolidating IT environment to rely on one vendor instead of multiple vendors

SOLUTION: Using IBM hardware and Linux to move from Sun Microsystems-based solution



ONLINE: Contact us and get more information about migrating from Solaris to Linux: www.ibm.com/linux

to make sure we could start small and then scale as the business grows,” Yu says. “Our IT environment allows us to do that.”

Jim Utsler, MSP TechMedia senior writer, has been covering the technology beat for nearly a decade. Jim can be reached at jutsler@msptechmedia.com.



Helping Solve Nature's

How Blue Gene/L is changing research and education
at Iowa State University **BY SHIRLEY S. SAVAGE**

Unlocking Mother Nature's mysteries isn't easy. Just ask any researcher involved in studying genomes and finding the genes that make up an individual organism. Figuring out a genome is like being given a huge box of jigsaw puzzle pieces where each big puzzle piece is made up of many small pieces. To solve the puzzle, you must fit the tiny pieces together to form larger pieces. Then, you've got to assemble the larger pieces into the proper structure. Completing the puzzle is a test of the researcher's patience and fortitude. But suppose the researcher had an able assistant that could process the little pieces into the larger pieces quickly and efficiently; you'd be able to see the overall structure faster. Additionally, you'd be able to spot those little discrepancies included in the genome—like repeated sequences. Such an able assistant does exist. It's called IBM* Blue Gene*/L. And it's helping researchers and students at Iowa State University solve one of the most complex biological genomes known.

In January 2006, Blue Gene/L arrived on the university's Ames, Iowa, campus and made an impact. Iowa State is one of five U.S. universities with a Blue Gene/L system. The acquisition of Blue Gene/L has catapulted Iowa State into the top 100 on the 2006 TOP500 list of supercomputers, with a worldwide ranking of 99. The 2,048 processor, 11 trillion-byte machine runs Linux* on a POWER* processor-based system.

"This is the first teraflop system we've ever had at Iowa State," says Srinivas Aluru, a professor in the electrical and computer engineering department and chair of the bioinformatics and computational biology program. "It's as much as 20 times more powerful than the next-most powerful machine we've had on campus. It has significantly enhanced our high-performance computing capabilities. In April, we received a best paper award from IEEE's (Institute of Electrical and Electronics Engineers) International Parallel and Distributed Computing Symposium for our work on maize assembly using Blue Gene/L. Without it, we wouldn't have been able to do the work."

The university acquired the supercomputer thanks to a National Science Foundation (NSF) grant and allocations from the Iowa State's President's Office; the Office of the Vice Provost for Research; Information Technology Services and the Plant Sciences Institute.

Once installed, Blue Gene/L, was immediately put to work. Aluru's group used the supercomputer to examine a large-scale problem in computational genomics known as expressed sequence tag (EST) clustering. The challenge in working with ESTs is to analyze the large datasets that comprise ESTs. Aluru's group chose to look at the mouse EST collection, which is maintained by the National Institutes of Health and has more than 3.7 million sequences. How long did it take to put that puzzle

Mysteries

together? “We were able to analyze them in less than 10 hours,” says Aluru. “So far, no one else has been able to demonstrate the ability to handle such a large collection directly. Most analyses are done on less than 1 million sequences.”

The Linux Advantage

One of the keys to the superior performance of Blue Gene/L is the Linux platform. “I’m a person who works on the algorithms and applications side of things, not the operating-system side. But that said, I can see the advantages of having a Linux implementation,” says Aluru. “If you look at the individual nodes, they need a very lightweight

being solved—the maize, or corn, genome. The maize genome is estimated to be as large as 2.5 billion nucleotides long, making it the third-largest genome to be sequenced, after the human and mouse genomes, which are 3 billion nucleotides long. Begun in November 2005, the maize project is led by Washington University in St. Louis, which performs sequencing. Iowa State is contributing the sequence processing and high-performance computing work on the project.

A few years ago, biologists did preparatory work for the project by sampling the entire maize genome using a technique called gene-enriched

sequencing gaps,” he says. The research team uses Blue Gene/L to do a combined analysis of all the BACs to find overlaps and fit the sequences together to build larger and larger ones.

Gaps aren’t the only challenges encountered by the biologists. Repetitive sequences are also an issue. The maize genome is full of repeats with a short evolutionary history. Some biologists believe the repeat content in the genome accounts for 60 to as much as 80 percent of the genome.

When there are repetitive sequences, researchers must make certain they aren’t mistaking the repeats as being from the same region. The sequences could be

UP CLOSE

CUSTOMER: Iowa State University

HEADQUARTERS: Ames, Iowa


BUSINESS: Higher education

HARDWARE: IBM BlueGene/L

SOFTWARE: Linux

CHALLENGE: Delivering fast and high-performance super computing to aid in cancer and genome research

SOLUTION: Running Linux and POWER processors on the BlueGene supercomputer



operating system derived from Linux. The rationale being that the machine is built for high performance. The goal isn’t to have a lot of people using the machine at the same time, but to deliver the highest performance. Linux, being open source and very adaptable, makes this easy to do.

“The individual nodes need to support the computation. When the computation is running, the nodes need to step away from all other tasks and let the application run. I don’t see how this can be done with a proprietary operating system or a very cumbersome operating system that does everything but can’t be tweaked in the right way.”

The Maize Maze

The speed and adaptability of Blue Gene/L is proving to be extremely helpful in one of the most complex puzzles

sequencing. In this technique, pieces of DNA are sampled from the genome where mostly genes reside. Starting in November 2005, biologists began to map the genome. Washington University scientists are breaking the genome into large chunks, known as bacterial artificial chromosomes (BACs), containing about 200,000 nucleotides. The BACs are then sequenced, with more than 1,200 BACs sequenced so far. Blue Gene/L is assisting with sorting the pieces. As the individual BACs are assembled, Iowa State is using Blue Gene/L to find out how the gene-enriched sequences fit into the individual BACs. That information will be used to perform better assemblies, observes Aluru.

However, assemblies don’t always go smoothly. “Sometimes you may not be able to get the whole sequence—you get it in two or three pieces with some

similar, but might actually be from another part of the genome. “The higher the repeat content, the greater chance of these mistakes,” Aluru says. That makes maize much trickier to sequence than the human or mouse genomes.

Confident the maize challenge will be solved, Aluru is already planning the next step. His team plans to compare the maize genome with rice (which has 450 million nucleotides) to see what genes are common and what genes are unique.

The maize project is just one of the applications for Blue Gene/L at Iowa State. It’s being used to support research in a number of areas, such as computational-biology tasks in comparative genomics and microdata analysis. “We are interested in going to public repositories, downloading thousands of experiments and collectively analyzing and discovering something

new that's hidden in the data," according to Aluru.

Blue Gene/L also is being used for plant-genome annotation to update a NSF-funded plant repository for protein structure prediction and refinement; quantum chromodynamics; particulate-cell simulation; nuclear physics; molecular dynamics and global-climate-system simulations.

Having Fun with Blue Gene/L

Although serious research is one of the major uses of Blue Gene/L, the supercomputer has a role to play with a more lighthearted educational task at the university. Each year, the Center on

faculty work with students to design programs and methods, but it's the students who most often are working on the system and building the codes. Having access to a teraflop computer attracts and keeps students on campus.

Blue Gene/L is being incorporated into the curriculum and will be used in the parallel-computing course taught each autumn. The course includes programming, developing applications on parallel systems and looking at developing scientific applications. Students will have access to both Blue Gene/L and a cluster, which will be used to discover programming bugs before running the programs on Blue Gene/L. "We'll be

when Patrick Schnable, a professor of agronomy at Iowa State and director of the Center for Plant Genomics, attended a conference where a few sorghum plant sequences were presented. Schnable asked Aluru to run the half million sorghum sequences available. "We were able to do the analysis and make it available on the same day using a Linux cluster," says Aluru. "When you can go back with this kind of analysis in a short time, I think the community will be more interested in learning about high-performance computing. Promoting the alliance between high-performance computing and biologists is very good



Information Assurance runs a cyber-defense competition open to high school and undergraduate students. In the competition, one group of students develops passwords that they believe are secure, while another group tries to crack them and now has the assistance of Blue Gene/L to do so. The process allows students to learn which passwords are vulnerable. "We blocked off the machine for a few days to be used in the cyber-defense competition, and it was a lot of fun," chuckles Aluru.

Students arriving from all over the country to participate in various programs such as summer school and science competitions are given a tour of the campus, which always includes a visit to Blue Gene/L. In fact, the predominate users of Blue Gene/L in research are actually students—graduates and post-doctoral candidates. The

teaching the theory of how large-scale parallel systems work. It's good for the students to actually get a feel of these systems. No matter how much theory you teach, unless the student can experience it, there's always something lacking," says Aluru.

The faculty is being encouraged to use this amazing resource. Aluru would like to see more engineering faculty—especially those working on computational fluid dynamics and molecular dynamics—use the supercomputer. He's trying to spread the word about high-performance computing among the biologists. "Many biologists aren't used to using much more than a Mac PC. I'm hoping to impact the community by consistently delivering research results that can't be done without this kind of machine," he says.

One proof of high performance computing's power occurred last year,

for the computing community. The integrated work between computer scientists and biologists is one of the top feats of the next century. We need to play our part in making this happen."

The Right Path

Whether it be finding a path through the complicated maze of the corn genome or opening up a world of possibilities to students, Aluru and Blue Gene/L are helping to make it happen. The university can be proud of the acquisition and for the lively advocacy of professor Aluru.

Shirley S. Savage, a Maine-based freelance writer, has written articles for IBM Systems Magazine. She also is the author of several management reports on technology and energy topics. Shirley can be reached at ronin@thinkingenergy.com.

The Seam



less Enterprise

Linux fits Success Apparel's business

BY JAMES MATHEWSON

Steve Golub knew he had to make a drastic change to his enterprise to fit it into a new, smaller datacenter. As the manager of operations for Success Apparel, a leading New York City-based children's apparel design and manufacturing firm serving Wal-Mart, Target, Kohl's, and J.C. Penny Co. Inc., and other retailers, Golub needed to consolidate 19 Microsoft* Windows* servers into seven machines to fit into the new datacenter. He decided to take the opportunity to create a more scalable, flexible and secure enterprise that better fits his company's needs. And the solution he chose not only fit the new datacenter, it accommodated two additional servers for redundancy and failover.

With the help of P&M Computers Inc., a Cliffside Park, N.J.-based enterprise-systems integrator, Golub chose Novell SUSE Linux* Enterprise Server running on IBM* System x* servers as the basis of the new datacenter. Other key components included Novell Open Enterprise Server, VMWare GSX virtualization software, IBM Tivoli* Storage Manager and Novell iPrint, which manages the company's huge printing demands.

With virtual Novell SUSE Linux Enterprise Server instances running on the new datacenter, the solution enables instances of IBM WebSphere* Host Integration, IBM Lotus* Domino* server and IBM Lotus Workplace* Services express. The combination offers both rich- and thin-client offerings for Success Apparel's widely distributed workforce.

"We have found that Novell Open Enterprise Server running Linux not only has greater integration, but makes for a much more secure and reliable network," Golub says.

"P&M Computers presented a great proposal and really steered us in the right direction," Golub adds. "The virtualization with Novell SUSE Linux Enterprise Server gives us tremendous cost and space savings."

Worldwide Collaboration

Success Apparel works with partners and distributors worldwide. Its employees rely on e-mail for most of their communication across multiple time zones. Its far-flung design team sends huge files that must be passed through the network and printed quickly. With so much international e-mail, the company's Windows infrastructure was vulnerable to viruses, and downtime was becoming an issue.

"At this point, there are not that many Linux viruses," Golub says. "Whether in the future there will be is anyone's guess. The important part as I see it, is I have much more control over what is and is not running on my Linux servers than I ever had with our previous or present Windows servers."

Golub says the company had some exposure to its Lotus Domino and Lotus Notes* e-mail solu-



UP CLOSE

CUSTOMER: Success Apparel

HEADQUARTERS: New York

BUSINESS: Children's apparel design and manufacturing firm

HARDWARE: IBM System x hardware, TotalStorage DS400

SOFTWARE: IBM Tivoli Storage Manager, WebSphere Host Integration, Lotus Domino server and Lotus Workplace Serves Express, Novel SUSE Linux Enterprise Server, Novell Open Enterprise Server, VMWare GSX virtualization software, Novell iPrint

CHALLENGE: Consolidating 19 Windows servers

SOLUTION: Using System x hardware to go from 19 Windows servers to seven System x servers and DS400s in order to consolidate and reduce hardware costs

tion in the past, but it was Windows technology-based. "As part of the adoption of Linux in the data-center we wanted an enterprise messaging system that ran on Linux and Domino fit the requirements perfectly," he says.

The company also improved e-mail performance and availability with Domino's clustering capabilities. And Domino gives Success Apparel's users hand-held capabilities they didn't have before, with a Domino Blackberry Enterprise Server.

"We are in a dynamic business with global implications in our supply chain, and many of our

human resources need to be connected seamlessly around the clock," Golub says. "It was vital that in a virtual environment that the e-mail and the Blackberry connection be rock-solid."

A Closet Full of Storage

Design work can be data-intensive. Designers must store large files, check them out to modify them and pass them around to colleagues for review. "With the help of P&M, we were able to locate a SAN solution that fit our budget needs and our virtualization needs," Golub says. "We choose the IBM TotalStorage* DS400 with additional IBM Fibre Channel (FC) switches."

Because we wanted offsite resiliency, our solution includes a switch and a DS400 in each location with a virtual clustered Domino Server and a few replicated file systems on each," Golub adds.

The DS400 is widely regarded as the most affordable storage solution for Intel* servers on the market. If more storage is needed, Golub can simply add additional DS400s to scale up without incurring large, incremental costs.

Stitching the Solution Together

A key way to squeeze a lot of computing power into a small space is to ensure that all of your resources are being utilized. For Success Apparel, this is done through virtualization with VMWare GSX software. Instead of running one application server per machine, as is common in Windows settings, Success Apparel can now host multiple application servers per Linux solution.

"At present, we still have an amazing amount of capacity to draw on for a new solution," Golub says. "VMWare's host-based virtualization allow us the flexibility to still run virtual server on a server that had proprietary hardware, as in our fax server as well as on our IBM Tivoli Storage Manager server. Both of which could not only be virtualized but also have workloads that allow for under-utilized capacity."

The increased utilization not only helps fit a lot of computing power in a small space, but it also allows Success Apparel to reduce its hardware costs, especially its maintenance and management costs. The results

its ROI target of less than 12 months for the new solution.

“We operate in an industry with small margins, so any savings from technology goes straight to the bottom line,” said Golub. “Because our new environment is so

But cost control isn’t the only benefit of virtualization. Golub also benefits from improved manageability of the integrated environment that virtualization brings. Using the Web-based interface of Novell iManager for centralized control, the IT staff can manage its entire network without physically visiting individual servers.

“With our previous Windows environment, it was harder for us to manage everything at once, like an orchestra without a conductor,” says Golub.

“Because our new environment is so much more stable and manageable than before, we have more time to spend finding new ways for technology to make a difference to our business.”

—Steve Golub, manager, operations, Success Apparel

include a reduction in hardware from 19 to nine servers, and reduced administration costs by an estimated 25 percent. That leaves more time for the existing IT staff to concentrate on new retail-technology solutions. Golub estimates that the company hit

much more stable and manageable than before, we have more time to spend finding new ways for technology to make a difference to our business.”



James Mathewson
is the technical knowledge offering manager for the IBM ISV Business Strategy and Enablement organization. He can be reached at jmath@us.ibm.com.

Past, Present and Future

**Three executives discuss
the whys and wherefores
of IBM's approach to
open source beyond Linux**

BY JIM UTSLER

The open-source community has come a long way since Linus Torvalds developed and then openly shared Linux*. Now, more and more applications, development tools and other technologies are being put up on the block, ready for developers of all stripes—from a guy working at home to an army of corporate programmers—to add to them.

Linux, of course, is still considered the flagship of all open-source projects, but others are making their way into mainstream corporate IT environments. *Linux Executive Report* (LER) recently discussed this trend with three IBM executives—Scott Handy, vice president, Worldwide Linux and Open Source; John Palfreyman, director of Open Source, Grid and Virtualization services; and Jeff Smith, vice president, Open Source and Linux Middleware—to get their perspectives.

LER: Why are customers now becoming more interested in open source beyond Linux?

JS: I think there's been quite a bit of success with Linux, and people are now saying, "Hey, I wonder if this can apply to other things." So one reason is the success of Linux, which leads users to look at other open-source options. Another thing is the new business model that surrounds open source, where you pay less for the license and more for the support and subscription. Those models are getting people's attention, and they're investigating whether the results of that can offer them new or better value than they've been able to get with whatever they have been using in the past. A third thing is that customers are realizing more and more the value of building solutions around open standards, and open source projects are a good way to get broad proliferation of open standards.

SH: That's right, because when they're using an open source technology, they're using something that everybody else is also using so they can get interoperability and adherence to a particular standard. Of course, most business are working in a mixed environment, with customers having both open source technology and maybe some private source as well.

JP: As Jeff points out, some customers consider open source because they're looking for cost savings; they're obviously attracted by the fact that the software does not attract any license fee. However, they should do so with both eyes open. Yes, license fees don't exist, but that doesn't mean there aren't costs involved in the project. So we need to make sure they consider all of the TCO (total cost of ownership) implications of what they're doing, and the other benefits that will accrue from making the right choices.

LER: What are some of the challenges customers are facing when adopting open source beyond Linux?

JP: Let's look at that a bit. We advise customers to look at their business

"Apache is important because it is an open-source community that proliferates the Web, open Web standards, and Web-applications standards."

—Jeff Smith, vice president,
Open Source and
Linux Middleware, IBM



needs first and not to adopt open source just for open source's sake. And the challenges here can be several fold. They need to determine how to create the appropriate combination of open-source and proprietary software, how to assess the maturity of the open-source software—how to make sure that the community support is there and it's not just from one vendor who's actually marketing their product in an open-source license—and to understand that there are certain things that need to be taken care of with open source, such as configuration management and control. With a traditional piece of software, the vendor will take care of this. This isn't the case with open source. Customers have to make sure they really understand the full and holistic implications of using open source in their business.

JS: Integration is also an issue. Customers who bring components of open source back into their environment find that they go through more work to integrate those things than they might with a vendor who built some of that integration into their commercial offerings, as in the case with the WebSphere* family. For example, WebSphere Application Server (WAS) uses the open source Apache HTTP server as a key component of an integrated software solution implementation.

SH: I think the challenges when implementing open source are the same as

with any technology deployment. Customers should have all of the same types of plans in place for supporting that technology that they would for any other private source technology. Once customers see it in that light, they're much more comfortable with the list of things they need to do to deal with in an open-source environment. They often can just use the same checklist that they've used when deploying any new technology, all the way down from how to get it approved on their standard deployment list to then rolling it out into production.

LER: Why is IBM so involved in the open-source space?

SH: Based on our experience with Linux, we figured out that some very useful things are developed with open source and that, often, it's actually in both the customers' and IBM's best interest to use open source in the model and the mechanism for how we deliver innovation to customers. From our perspective, it's certainly better than having multiple vendors doing duplicate development work to create more or less the same product. It's a waste of resources. But by sharing the exact same code base, we, as a vendor, can take our valuable resources and use them for developing value add on top of that standard, whether that be Linux, whether that be an open-source J2EE app server, whether that be Eclipse. This philosophy applies to every open-source project that we've backed. We can share the development work for the open source and open standards and then provide more value to the customer by implementing a unique value-add on top of it.

JP: We're also seeing customer interest in open source. And we see a market there because customers see competitive advantage if they are able to combine open source and private source, and we believe that IBM is well positioned to help them make objective, business-led decisions to maximize their competitive differentiation from harnessing this combination.

JS: In keeping with that thought, we've found that commercial ecosystems form around cores of open-source community innovation that include the business opportunities that John mentions. It's one of the reasons that open-source projects have gotten the attention and the investment that they have. Linux is a prime example of that. People are selling billions of dollars of hardware under Linux and middleware on top of Linux and services surrounding Linux. The starting premise is that there's a core of community-driven innovation that an ecosystem builds around. IBM's goals regarding that innovation are to participate in the ecosystem in a way that is mutually beneficial to our customers, to IBM and to the ecosystem itself.

"What you want to do is use the capabilities of open source where appropriate, which probably will be in a mixed open-source/private-source environment. If you can get a little in, get a little in, if you can get a lot in—the more the merrier."

—Scott Handy, vice president,
Worldwide Linux and
Open Source, IBM



LER: How has IBM's experience with Linux helped it with other open-source projects?

JS: First of all, we've learned about what does and doesn't work when you're working with open-source communities. We've become knowledgeable about the communities themselves: how they function, how to contribute to them and how to earn their respect,

which is important for you to be able to get value from them. We have shown that this idea of commercial ecosystems forming around these open-source community innovations can be successful, which has led us to consider more of them. And we've shown that customers are interested in this stuff, that alternative solutions, especially those that offer more choice to do things, are inherently interesting to customers. So not only do we see value from our business strategy perspective, but our customers to a certain degree are also voting for us to participate in these things.

SH: We've also learned that to make an open-source project successful, it has to be repeatable. Let's take a quick look at Linux: We learned that there was a certain set of things we knew we had to do to be successful in the community, including becoming a major contributor, not only for credibility with the customers, but also for credibility within the community. And that can only be done through real technical contributions of innovation into the open source community. That allows us to advance the technology so that it meets the needs of our customers, so if we have customer requirements, we know how to feed those requirements into the technology. It also allows us to provide support offerings for the technology because we're familiar with the code base and have a mechanism to support our customers.

JP: I'd add to that. Because we're extending our services portfolio to embrace open source beyond Linux, we're extending the skill sets of our professionals who already have experience with Linux. We're also used to interacting with a product that is open source, so this gives us a sound knowledge base to build on. However, it would be foolish of us to expect that the entire spectrum of open source will evolve in exactly the same way as Linux. Open source is very multifaceted, with penetration into areas such as tooling and middleware and embryonic penetration into the application arena.

LER: What are some of the major areas of IBM investment in open source?

JP: Our main areas of investment largely align with what the Software Group has been doing. We're clearly looking to make sure that we include open source in our services portfolio, to build on the investment in proactive community development that has been made in the Software Group, with things like Eclipse and Apache and the other communities that we're very involved in.

SH: Right now, there are several areas that I'm focused on. One involves the application-server space, which we've done by adopting Apache as our default HTTP server in WebSphere and continue to do by working within the ecosystem development around Apache Geronimo. We also created the Eclipse project when we open sourced our technology, which involved 40 million dollars worth of development, back in November 2001. Since then, Eclipse has been a phenomenally successful project, with more than 800 tools developed for it. And building on Eclipse in the area of the rich-client platform (RCP) that was in Eclipse 3.0. That was really for extending Eclipse beyond app-development tools to support building client-side apps. So instead of developing tools, we're developing apps. This is a very important open-source based project for us because this really is nothing short of re-creating what we did on the server, of having a single programming model that is multi-OS and based on open standards to provide customers with a more cost-effective way to deploy apps that are multi-vendor and multi-OS. We also felt that there were a lot of benefits to using middleware on the server, so we injected client-side middleware into the strategy and we did that by developing our own framework on top of Eclipse 3.0, which is the IBM Workplace* managed client.

We noticed that customers had to deal with storage management that was vendor unique in all cases. For example, IBM was solving it one way with our

storage, EMC another way—all the various data-storage vendors had their unique way of doing storage management. All customers wanted was to have a common storage-management solution. So IBM took a leadership position and open sourced some of our technology as well as contributions from other companies to launch the Aperi project, which is building an open-source implementation of the open standard SMI (Storage Management Initiative) backed by SNIA. And moving beyond software, we then started thinking we could extend the same open-source software concept to hardware. We're now licensing the specifications for the Power architecture* to interested companies, and there are currently around 40 companies that have joined power.org. That's one hardware implementation. Another one we did was with blade.org. So we're already sharing the specifications for the blade and chassis architecture with other companies, and now there are more than 600 companies that are involved in blade.org. Another area that I'm focused on involves grid. We started getting involved in open source grid project back in the year 2000, and there was a whole specification that was being created called open grid services architecture—OGSA—and we've implemented that grid architecture specification into our own products, such as WebSphere and Tivoli* software.

JS: We also still have hundreds of people writing code for Linux fulltime, contributing into the Linux open-source community, porting Linux to IBM hardware platforms or writing middleware for it, so the open-source involvement with IBM on Linux is fundamentally the biggest. The second biggest is probably Apache. I don't want to repeat too much of what Scott has already said, but Apache is important because it is an open-source community that proliferates the Web, open Web standards, and Web-applications standards and Web-services standards that are fundamental to our software strategy. As an example, we announced

"Because we're extending our services portfolio to embrace open source beyond Linux, we're extending the skill sets of our professionals who already have experience with Linux."

—John Palfreyman,
director of Open Source,
Grid and Virtualization services



the availability late last fall of WebSphere Application Server Community Edition (WAS CE), which is based on an open source J2EE application server project from Apache called Geronimo, and we even bought a company that had a commercial offering based on this technology that formed the basis of our WAS-CE offering. We donated some embedded Java* database technology a couple of years ago, and that has now graduated to a full-blown database project inside Apache called Derby. We continue to offer that as a commercial offering called CloudScape as well as a bundled version with the WAS CE offering. And then, of course, there's Eclipse, which may be just as big as Apache ... at least from an IBM participation perspective.

JP: The other important thing to stress is that we decided to join up on our open-source plans to make sure that there are no overlaps and gaps in our approach. So that's something that's really exciting—actually having something truly coordinated between the three different divisions—software, systems and services—and getting a lot of benefit from bouncing things off each other and moving these activities forward together.


LER: What advice would you give to IBM customers considering adopting open source?

JS: First of all, if you're used to buying software in a traditional way, where you buy it up front from a particular vendor and it's their own technology, then there's never been any question about where you're going to get support or how much support you're going to get. But what customers are finding with open-source projects is that there's no guarantee there's any support at all; that you actually have to pick where you're going to get the support from, because it's often offered by more than one vendor. Or it could be that the only support available is the Web interface into the community, in which case it's not going to be as predictable as you might like if you're an enterprise customer. Support options may exist, but you have to deliberately seek them out.

JP: The advice I would give is stand back and consider your real business rationale before acting. Open source must never be viewed as the universal answer to all questions, but if used in the right combination with traditional software it can yield significant business benefit, allowing the customer to innovate. If the answer is "We don't need open source yet because it's really not mature enough in the area that we want to do business," that's a really good answer. Increasingly, however, that won't be the answer.

SH: In keeping with that, what you want to do is use the capabilities of open source where appropriate, which probably will be in a mixed open-source/private-source environment. If you can get a little open source into a project, get a little in, if you can get a lot in—the more the merrier. But don't set these arbitrary rules that everything has to be 100-percent open source when 30 percent might be enough to get some advantage. Take your success and expand on it. In short, just bite off what you can chew and go from there.

Jim Utsler, MSP TechMedia senior writer, has been covering the technology beat for nearly a decade. Jim can be reached at jutsler@msptechmedia.com.



Tony Parziale, CIO, Palm Beach Community College, chose IBM hardware and Linux to help consolidate and virtualize the school's IT environment.

**How Palm Beach
Community College learned
that consolidation and
virtualization can result in more
than just ease of administration
and lower overhead costs**

BY JIM UTSLER

reating a community

Most often, the primary reasons for undertaking a consolidation effort are related to the bottom line: increased ease of management and administration, and reduced costs and datacenter space. After all, dealing with multiple stand-alone servers can be a chore, with their varying operating systems (OSs) and applications, direct-attached storage and expensive software-licensing schemes.

Given all that, it's really not surprising that more organizations are turning toward consolidation to ease their IT woes. But in the frenzy to do so, they may be overlooking other benefits that can arise from their efforts. For example, consolidation allows organizations to more closely integrate their systems to offer better services, such as portal access to personal information.

So as organizations begin considering or even implementing consolidation efforts, they should reflect on what real-world benefits can be gained.

As Palm Beach Community College recently learned, consolidation isn't so much about the nuts and bolts of a technological deployment, but how that deployment is going to benefit end users.

Tony Parziale, CIO of Palm Beach Community College, notes, "We have five locations scattered

about the county. By consolidating student records on compatible systems, the students can access their data—such as financial-aid or transcript information—from any of these locations. In the past, we would have had to print out the information and fax it to the other campuses."

In the case of Palm Beach Community College, the value of technological consolidation and virtualization came in the form of a tightly integrated mix of a



UP CLOSE

CUSTOMER:

Palm Beach Community College

HEADQUARTERS: Lake Worth, Fla.

BUSINESS: Higher education

HARDWARE: An IBM System z z890, two IBM BladeCenter servers, an IBM System Storage DS6800 storage device and an IBM System Storage 3583 Ultrium Scalable Tape Library

SOFTWARE: SAN Volume Controller, Software AG Adabas database-management system, Tivoli Storage Manager, SUSE Linux

CHALLENGE: Consolidating and virtualizing its IT environment

SOLUTION: Running several Linux partitions on the z890, consolidating many Windows technology-based servers to the BladeCenter servers and using IBM's SVC to virtualize its storage environment

new IBM* System z* 890, two IBM BladeCenter* systems, an IBM System Storage* DS6800 device and IBM virtualization technologies, such as the IBM SAN Volume Controller (SVC). And so far, this newly revamped consolidation and virtualization effort has received nothing but straight As.

Leading the Pack

Established in 1933, Palm Beach Community College was Florida's

first public community college. As its name implies, it serves Florida's Palm Beach County, which Parziale says "is bigger in land area than Delaware or Rhode Island." It has five locations in Lake Worth, Belle Glade, Boca Raton, Palm Beach Gardens and West Palm Beach, and about 49,000 enrolled students who are eligible to take part in more than 90 degree and certificate programs.

As one might expect, an organization of this magnitude, with its many processor-intensive transactions—according to Parziale, more than 85 percent of its students register online—requires a heavy-duty IT environment to support its users, who include not only students, but also faculty and staff. To that end, it currently has a System z890 platform hosting five SUSE Linux* partitions, two fully equipped BladeCenter servers, a DS6800 with 10 TB of storage and an IBM System Storage 3583 Ultrium Scalable Tape Library. It also runs a proprietary ERP application cooperatively developed by seven Florida schools that are part of what Parziale calls "the Florida Community College Software Consortium." These schools include Broward

Community College, Indian River Community College, Miami-Dade College, Tallahassee Community College, Okaloosa-Walton Community College, Polk Community college and Palm Beach Community College, as well as associate members Mercer County Community College in New Jersey and Louisiana State University at Shreveport.

This ERP application, which was built on the Software AG Adabas data-

base-management system running in the z890 server's Linux* environment, handles the school's student system. According to Parziale, this student system "consists of everything from registration, financial aid, and certain types of billing, financials and HR." This consortium notion of building an application allows each school to add functionality to it and then share that functionality with the other schools, in a development method not unlike that of the open-source community.

"Each of the schools has its own programming staff, so what we do is build functionality and then share it back with the rest of the consortium. Central to this is a consortium office with its own IT staff that we have located in Fort Lauderdale, which combines this functionality and releases it to all of the schools on the next release of the application," Parziale explains. "The fees you pay to the consortium are based on your enrollment numbers."

Because of this cooperative agreement, many of the other schools involved in this group, including Broward Community College and Indian River Community College, are also moving to the z890 platform as a way to create similarly outfitted IT environments. In fact, all of the schools are moving to Linux and UNIX* to create a level application playing field.

"The value of this consortium relationship is that one school can invest a lot of effort into moving forward either regarding applications or platforms and then share that knowledge with the other schools. We just happen to be an early adopter in this area," Parziale notes.

As part of that leadership, the school decided to upgrade its older IBM mainframe to a newer one capable of running Linux, which would provide more flexi-

bility regarding application and server choice should it decide to migrate to another platform. One of the reasons for the urgency to upgrade was that the consortium's licensing agreement for Software AG for z/OS* was due to expire at the end of 2006. Consortium members decided to approach Software AG to determine if it would support Adabas for Linux on an IBM System p5* platform (at the same time, it was reviewing other UNIX* technology-based offerings from other vendors). Software AG's response was that it wouldn't support Adabas for Linux on the System p5 platform (only AIX*), but that it would for Linux on the System z or an Intel* technology-based platform. "That limited our choices," Parziale remarks.

As a result, the consortium approached IBM about running Linux on the System z platform, not necessarily wanting to move to a Windows* technology-based platform, already having too many Windows servers operating in the schools' various IT environments. Based on a number of factors, including the platform's flexibility, scalability and "the breadth of development behind the System z platform and the tools that are available for it," Parziale says, the consortium (members of which were operating in a

mix of z/OS and Virtual Storage Extended (VSE) environments) decided to move to the System z platform (specifically, the z890 servers).

"We wanted a platform that could handle any direction the consortium decides to go in in the future, whether it was a prepackaged application or another database," Parziale continues. The school's z890 platform now has five virtualized Linux partitions, most of which are dedicated to the ERP system, including for production, development and testing.

Continuing Opportunities

The decision to move to the System z platform made, Palm Beach Community College also decided to introduce the BladeCenter servers into its IT environment. This would help reduce its reliance on the 70 or so stand-alone servers it had been using for a variety of purposes, including distance learning.

"Each of them had its own applications and storage, so you have to manage each one individually," Parziale says. Now, about half of those stand-alone servers have been migrated to the BladeCenter servers populated with HS20 blades, and the school is using VMWare to virtualize those servers for additional consolidation purposes. Using these virtualized servers, the organization can also pilot applications without having to buy the hardware to support them, a huge cost savings.

At the same time, in July 2005, Palm Beach Community College also brought in the DS6800. As part of its overall consolidation effort, it wanted a single storage platform, no matter where the data came from—be it the System z890 or its BladeCenter server. To make the most of this move, the school deployed IBM's SVC, which would allow it to

"It's not all about the technology and saving money and time, but enhancing everything Palm Beach Community College has to offer."

**— Tony Parziale, CIO,
Palm Beach Community College**

virtualize storage across all its platforms. Rather than having direct-attached storage that may have been underutilized, it now has a single virtualized storage pool that has a much better utilization rate.

“We might have had a server with 200 GB of storage but were only using 20,” notes Parziale. “It was a waste of resources.”

Backing up this system is the IBM 3583 Ultrium Scalable Tape Library. In the past, backups were a somewhat muddled affair, with the mainframe-generated data and Windows technology-based data having to be backed up to different devices. And in the case of that environment, backups were taking up to 24 hours to complete.

“Now,” says Parziale, “we put the four-drive 3583 unit in place and we’re backing up off the SAN with Tivoli* Storage Manager (TSM). We’ve been able to cut the backups for our entire enterprise down to five hours.”

‘Single Cockpit View’

Although this consortium is an important reason why Palm Beach Community College decided to consolidate many of its IT assets, the school had additional reasons as well. For example, it wants to run WebSphere* on Linux on the System z platform to power a college-wide portal system. Although consultants had suggested running WebSphere on either a Windows technology-based platform or as part of z/OS, the school insisted that it run in a virtualized Linux partition on the mainframe. In February of this year, IBM and the college ran a proof of concept regarding exactly that.

Once operational, this portal will allow students to have a single point of entry for all of the college’s services, including e-mail, grades, class registration,

messaging and course management. As Parziale describes it, “It would be a single cockpit view of their Palm Beach Community College life.” The school also wants to launch similar portals for staff and faculty, all of them using WebSphere running within Linux on the z890 platform.

“We continue to look for new opportunities to leverage the System z890 platform whenever we can,” remarks Parziale. This includes moving its distance-learning application and help-desk system to Linux on the System z platform. Parziale is also working with IBM to move its Microsoft* Exchange-type applications to the z890 while retaining the Outlook presentation on the client. “We’re not afraid of attempting anything at this point,” he says.

“We’re able to personalize the online college experience for our users so they can appreciate the entire learning process. And we’re proud of that. After all, that’s what we’re here for. It’s not all about the technology and saving money and time, but enhancing everything Palm Beach Community College has to offer.”

— Tony Parziale

This includes using its new virtualized-storage environment to image many of its paper-based documents. This will allow the school to not only safeguard, for example, student records—which, given its hurricane-prone location, is a necessity, as Hurricane Katrina proved—but also make them available online, no matter where a student might be. And as documents are imaged, the paper copy will be destroyed, freeing up valuable floor space and negating the need for the purchase of more filing cabinets.

Something to Be Proud Of

Although Palm Beach Community College had sound reasons for wanting to consolidate its IT environment (a process that took four or five months to complete, including a month of installation and testing), it has since discovered that consolidation provides benefits beyond easing administrative overhead and reducing hard costs, such as on new servers and storage. It also realized that a combination of consolidation and virtualization can create many new opportunities, including efforts such as the school’s portal and imaging initiatives.

Referring to the student portal, Parziale sums it up by saying, “We’re able to personalize the online college experience for our users so they can appreciate the entire learning process. And we’re proud of that. After all, that’s what we’re here for. It’s not all about the technology and saving money and time, but enhancing everything Palm Beach Community College has to offer.”

Jim Utsler, *MSP TechMedia senior writer*, has been covering the technology beat for nearly a decade. Jim can be reached at jutsler@msptechmedia.com.



Quotable Quotes

Customers explain the business benefits of deploying Linux

Aitana SBS

“On [the System i* platform] you can manage the resources and security from OS/400*, and run a high-speed virtual LAN between partitions with no new skill set required, which represents an immediate management-cost benefit. We can see [the System i platform] developing into one of the best platforms to support Java* applications running in WebSphere*, and we will be encouraging our customers to move away from RPG to become more open, and more modern. [The System i platform] is an excellent environment for serious, business-critical applications, and running Linux* makes it a great platform for the future.”

—José Carlos Ramos, e-business Manager, Valencia, Spain

Hennepin County

“It was a plus that IBM* WebSphere* would work seamlessly with the infrastructure we already had in place. With the consolidation of applications onto virtualized mainframe servers, the potential for county-wide savings is in place. We have spent an immense amount of redundant staff time and money licensing, patching and upgrading all our Microsoft* Windows* servers. The decision to run WebSphere Application Server on Linux* on [the System z* platform] has given us a flexible, enterprise-quality environment that helps free developers to write applications instead of supporting server infrastructure. Operations can create a new server for us within minutes—without requiring us to purchase more hardware and software.”

—Tammi Kolasa, information technology supervisor, Minnesota

Mobil Travel Guide

“IBM has optimized WebSphere* software for Linux*. Combined, they provide the high performance and low cost that our online service requires to respond to the needs of our customers.”

—Paul Mercurio,
senior vice president and CIO,
Park Ridge, Ill.

Conven Corporation Limited

“IBM* WebSphere* and Linux* are delivering good performance for banking-transaction processing, being a one-stop purchase from IBM* hardware and software not only satisfying our reliability and scalability requirement, but also with a major selling point—offering the lowest total cost of ownership.”

—Dr. Rudy Law, director,
Hong Kong

Conquering Open-Source Fears

BY SHIRLEY S. SAVAGE

ALL TOO OFTEN, companies greet the idea of open-source computing as a threat to the way business is currently done. One of the challenges faced by open-source advocates is to change that perspective. Open-source computing will change business as we know it. But rather than viewing it as a threat, it's helpful to see open-source computing as a business opportunity that needs to be embraced.

Linux Executive Report (LER) recently spoke with Dr. Robert S. Sutor, vice president, Standards and Open Source at IBM, to gain some insight on how to accomplish that change of view successfully.

(If you'd like to read more about Sutor's thoughts on life, open standards, and open source, visit: <http://sutor.com/blog>.)

LER: Why is it so difficult for companies to adjust to open-source computing?

RS: It's not unreasonable for people to have a natural resistance to open-sourcing software, particularly if they grew up in a culture that was about developing and then licensing software for money. Therefore, executives need to be much more explicit about open-source goals and laying down a number of steps to reach that goal. One of those steps is a business plan. If a company is trying to move to an open-sourcing world, what are the business goals? What will the company look like in terms of its products, the people to

help make this transition, and the revenue structure?

Executives need to ask: Is my company currently equipped personnel-wise to make the transition to open sourcing? I'm not suggesting that you fire everyone. Rather, there needs to be training, discussion, and very open

and participating in a strong open-source community can lead to a lot of innovation, which will help up create those connections and build things that people haven't previously envisioned. When innovation happens, there will be a lot of growth opportunities for those who choose seize them.



"It's not unreasonable for people to have a natural resistance to open sourcing software, particularly if they grew up in a culture that was about developing and then licensing software for money."

—Dr. Robert S. Sutor, vice president, Standards and Open Source, IBM

communication since people can harbor deep-seated reservations about moving to open sourcing.

LER: Isn't it all about conquering the fear of the unknown?

RS: In the last 150 years or so, people have gotten used to the concept of ideas being discovered and then owned. If you go back farther in history to everyone's favorite inventor and master, Leonardo da Vinci, the goal in making discoveries was to advance the general knowledge of civilization. Much of this knowledge created hundreds of years ago was a rediscovery of what people once knew as well as an incredible burst of creativity. In da Vinci's time, there was a lot of free sharing of ideas, which enabled people to make new connections between them.

IBM very much believes that nurturing

LER: What is your advice to those who are reluctant to embrace open-source computing?

RS: View open sourcing as a way for brand-new markets to be created, to have new product categories, and to have new ways to serve your customers. Those new ways might not be traditional, but could be profitable.

We're living in an age now that's almost a Renaissance period for software. People are more willing to experiment with the technology and the business models. We don't know where it's going. But at least we are creating the environment in which wonderful things can happen.

Shirley S. Savage, a Maine-based freelance writer, has written articles for *IBM Systems Magazine*. Shirley can be reached at ssavage@maine.rr.com.