



Forum:
How Open Is 'Open'?
Industry luminaries join the debate

Computing:
Multi-Core Debugging
and Performance
Enhancement



ENTERPRISE OpenSource MAGAZINE

Visit us at opensource.sys-con.com

THE LEADING MAGAZINE FOR ENTERPRISE AND IT MANAGEMENT

JANUARY 2007 VOLUME 5 ISSUE 1

SECURITY: ELIMINATING INSIDER THREATS

AS AN ESSENTIAL
SECURITY AND
COMPLIANCE
BEST-PRACTICES
MEASURE

PLUS

THE GENESIS OF THE
LINUX FOUNDATION

BITLEAP OFFSITE BACKUP

NEXT-GENERATION
LINUX CLUSTERS

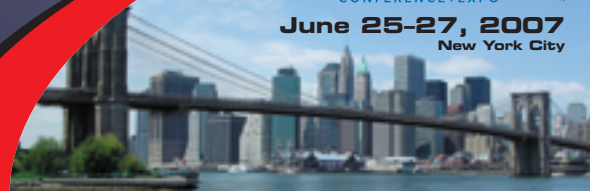


PRESORTED
STANDARD
US POSTAGE
PAID
ST. CROIX PRESS

End Annual

ENTERPRISE > 2007
OPENSOURCE
CONFERENCE+EXPO

June 25-27, 2007
New York City



SEE PAGE 27

IT'S THE END OF THE LINE FOR SPAM.



PUT IT IN THE CAN.

Ads, fraud schemes, porn, viruses, phishing attempts...
if it's offending e-mail, Roaring Penguin Software's CanIt™
solutions will stop it before it reaches your mail server.



- Simple administration
- Automatic spam and virus updates
- Compatible with all e-mail servers
- Available as software or hardware appliance
- Flexible end-user controls

Throw the book at spam today!
Request a free 20-day CanIt evaluation
and stop spam on YOUR terms!

ROARING 
PENGUIN
SOFTWARE INC.

What's Inside...

JANUARY 2007 VOL.5 ISSUE 1

**[5] INDUSTRY COMMENTARY:
The Genesis of the
Linux Foundation**

by Jim Zemlin

**[8] CASE STUDY:
BitLeap Offsite
Backup***How open source software powers
mission-critical backup services*

by Guy Suter and Ian Berry

**[12] ARCHITECTURE:
Next-Generation
Linux Clusters***An innovative architecture for
maximizing return on investment*

by Bob Monkman

**[18] SECURITY:
Trust But Verify***Eliminating insider threats as an
essential security and compliance**best-practices measure*

by Ellen Libenson

**[22] COMPUTING:
Multi-Core Debugging
and Performance
Enhancement***Additional pressures on
complex applications*

by Douglas Eadline and Vince Hauber

**[26] FORUM:
How Open Is 'Open'?***Industry luminaries join the debate*

by Jeremy Geelan

**[30] NEWS:
Open Source News**When Open
Is Really 'Open'

By Mark R. Hinkle

Of late there has been a lot of buzz around what constitutes open source software. Industry leaders and pundits alike have weighed in on the “openness” of certain software and companies’ business models. The generally recognized test for open source software is the Open Source Definition drafted by Bruce Perens and endorsed by the Open Source Initiative, the non-profit organization that shepherds open source licenses and gives an industry-recognized stamp of approval for these licenses.

The Open Source Definition includes 10 criteria that define what constitutes an open source piece of software. Recently, a number of companies have launched products that meet a limited number of the criteria, while still waving the open source banner. The issue is being brought to light as more users are adopting open source software, and more software vendors are adopting open source business models. This has many established community members up in arms as they see these companies as having a lesser commitment to the ideals of open source software and providing a disservice to the open source movement. The question is: “What can be done about it, if anything?”

The Open Source Initiative’s power stems mostly from public favor; they don’t own open source as a trademark. Their Website notes the rising ambiguity and offers the following advice:

While there is agreement on the broad term “open source” as meaning approximately what is captured in the Open Source Definition the term has, ironically, now become so popular that it has lost some of its precision. We strongly encourage everyone who cares about open software to use the term only to describe licenses that conform to the OSD, or software distributed under such licenses; but since the term has passed into more general use, we also encourage people to refer to the “OSI Certified” mark, which has precision and legal force in identifying software distributed under licenses that are known to meet the OSD requirements.

Those intimate with open source generally know what the spirit and intent of open source is, and acknowledge the Open Source Initiative as the “governing body” for blessing open source licenses. The real fear should be that end users are going to accept these *pseudo* open source products as having the same benefits as *real* open source software. Consequently, if, or perhaps when, they have a bad result they will consider this to be the state of open source software in general.

You may be asking yourself, as an end user of software, should you care? For many of us the point of contention is the ability of the software to meet your needs at a cost that’s acceptable. Cost is more than just dollars and cents. It may be the cost of the manpower to deploy and maintain the software. It could be the cost for vendor support contracts. It might be the cost to integrate a piece of software with another system. The cost of a piece of software relative to the benefit it provides is value. In many cases, the cost is so low it provides a much greater value than proprietary software. One hidden cost may be the lack of freedom you concede by not using software that is licensed under an OSI-approved license.

Most people probably lack any real religion when it comes to the open source movement. However, no matter whether you are using open or proprietary software, you should be aware of



—continued on page 7

About the Author

Mark R. Hinkle, editor-in-chief of *Enterprise Open Source Magazine*, is the vice president, Community and Business Development at Zenoss Inc. He serves as a founder of the Open Management consortium and is the author “Windows to Linux Business Desktop Migration” (Charles River Media).

mrhinkle@encoreopus.com



YOU ALWAYS HAD THE BRAINS. IT WAS THE TECHNOLOGY THAT WAS A LITTLE SCATTERED.

The HP BladeSystem c-Class with Insight Control Management.

The intuitive HP BladeSystem c-Class thinks just like you do — letting you monitor your infrastructure while helping to analyze your future needs. First, HP's OnBoard Administrator gives you out-of-the-box setup and configuration combined with power, cooling and enclosure management. After that, the Insight Control software steps in to let you control the rest of your environment, locally or remotely. And thanks to the integrated Insight Display — our unique LCD screen — you can interact right at the source to manage, deploy or troubleshoot.

Simply plug in the HP ProLiant BL460c server blade, featuring Dual-Core Intel® Xeon® Processors, and you'll get faster performance and versatility to support 32- and 64-bit computing environments. Use the HP BladeSystem c-Class for your business and you'll experience greater control over your time and resources.



Experience the HP BladeSystem and download the IDC White Paper "Enabling Technologies for Blade Management."



Click YouAlwaysHadIt.com/brains7
Call 1-866-625-4085
Visit your local reseller



Dual-Core is a new technology designed to improve performance of multithreaded software products and hardware-aware multitasking operating systems and may require appropriate operating system software for full benefit; check with software provider to determine suitability; not all customers or software applications will necessarily benefit from use of this technology. Requires a separately purchased 64-bit operating system and 64-bit software products to take advantage of the 64-bit processing capabilities of the Dual-Core Intel Xeon Processor. Given the wide range of software applications available, performance of a system including a 64-bit operating system will vary. Intel's numbering is not a measurement of higher performance. Intel, the Intel Logo, Xeon and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. The information contained herein is subject to change without notice. ©2006 Hewlett-Packard Development Company, L.P.

EDITORIAL BOARD

Editor-in-Chief

Mark R. Hinkle mrhinkle@encoreopus.com

Contributing Editor

Ibrahim Haddad ibrahim.haddad@sys-con.com

Migration Editor

Jon Walker jwalker@sys-con.com

Commercialization Editor

Paul Sterne sterner@sys-con.com

Desktop Technology Editor

Tim Griffin tim@sys-con.com

Review Editor

Matt Frye mattfrye@sys-con.com

Editor

Philip Peake philip.peake@sys-con.com

Contributing Editors

Kevin Larue kevinl@linspireinc.com

Christopher Negus cnegus@mwmt.net

Contributor

Rob Jones rob@hotlinuxjobs.com

INTERNATIONAL ADVISORY BOARD

Wim Coekaerts; *Director of Linux Engineering, Oracle*

Brian E. Ferguson; *Partner, McDermott, Will & Emery*

John Fowler; *Executive VP, Network Systems Group, Sun Microsystems*

Gaël Duval; *Cofounder/Director of Communication, MandrakeSoft*

Samuel J. Greenblatt; *Sr VP and Chief Architect, Linux Tech. Group CA*

Scott Handy; *VP, Linux Strategy and Market Development, IBM*

Bruce Perens; *Perens, LLC*

Stacey Quandt; *Principal Analyst, Quandt Analytics*

Thomas Reardon; *VP and GM, Client Product Group Openwave Systems*

John Weathersby; *Executive Director, Open Source Software Institute*

Ranjit Nevatia; *Director of Linux Strategy, VERITAS*

Andy Astor; *co-founder and CEO, EnterpriseDB*

EDITORIAL

Executive Editor

Nancy Valentine nancy@sys-con.com

Research Editor

Bahadır Karuv, PhD bahadir@sys-con.com

OFFICES

SYS-CON MEDIA

577 Chestnut Ridge Rd. • Woodcliff Lake, NJ 07677

Telephone: 201 802-3000 • Fax: 201 782-9600

Enterprise Open Source Magazine (ISSN #PENDING)

is published monthly (12 times a year)

by SYS-CON Publications, Inc.

Postmaster send address changes to:

ENTERPRISE OPEN SOURCE MAGAZINE

SYS-CON MEDIA • 577 Chestnut Ridge Rd. • Woodcliff Lake, NJ 07677

Copyright © 2007 by SYS-CON Publications, Inc.

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy or any information, storage and retrieval system, without written permission. For promotional reprints, contact reprint coordinator.SYS-CON Publications, Inc., reserves the right to revise, republish and authorize its readers to use the articles submitted for publication. All brand and product names used on these pages are trade names, service marks, or trademarks of their respective companies.

WorldWide Newsstand Distribution

Curtis Circulation Company, New Milford, NJ

FOR LIST RENTAL INFORMATION:

Kevin Collopy: 845 731-2684, kevin.collopy@edithroman.com

Frank Cipolla: 845 731-3832, frank.cipolla@epostdirect.com

The Genesis of the Linux Foundation

by Jim Zemlin

On January 21, the *New York Times* published a story on Linux. This wasn't an article on technical advancement: no new kernel or distribution had been released. It wasn't financial; there wasn't yet another impressive quarter from one of the many companies that build their business around Linux. Thankfully, it wasn't another piece of FUD about open source legal issues and dubious patent assertions from desperate competitors. Instead the article simply stated: "The Linux industry has united to compete against proprietary platforms." The Linux Foundation was born.

Late last year, our members and internal management teams decided the time was right to merge the two leading Linux consortia: the Free Standards Group (FSG) and the Open Source Development Labs (OSDL). Why now?

Since OSDL and the FSG were formed, more than six years ago, Linux has grown enormously in server, desktop, and embedded usage around the world—driving more than \$15 billion in annual sales of hardware, software, and services according to market research firm Gartner Group. Moreover, the open source model now dominates new software development and provides faster demand-side learning, higher quality, better security, shorter development cycles, and lower prices. OSDL and the FSG were helpful in this phase.

Now that Linux has achieved widespread use, it faces a different set of challenges. In order to advance the platform further, the Linux industry formed the Linux Foundation to standardize, protect, and promote the Linux platform. The distributed development and sales and support model of Linux drives many of its benefits, yet also creates challenges that can hinder its success. We want Linux to continue to challenge the dominant operating system in the world. We want Linux to offer true choice for end users everywhere, regardless of economic means or technical literacy. In order to do that, we need to keep the freedoms and advantages of the open source model while continuing to improve the platform and its competitiveness.

The Linux Foundation will work with our members to provide services that an open source platform needs to compete. These projects can be far-ranging, such as standardizing Linux so application developers can more easily target the platform, or smaller in scope, such as providing the legal infrastructure so open source developers can sign required NDAs before writing device drivers for private companies. Our projects can be technically complex, such as the new LSB testing framework that links compatibility tests to code development, or simple and straightforward, such as providing a neutral voice of reason to the press to counter competitors' aggressive PR tactics.

The Linux Foundation has united the Linux ecosystem with founding platinum members Fujitsu, Hitachi, HP, IBM, Intel, NEC, Novell, and Oracle. Other members include AMD, Cisco, Dell, EMC, Google, Red Hat, Sun Microsystems and more than 70 other companies, industry end users, universities, and community groups. We have representation on the board from the Technical Advisory Board and key Linux individuals so the technical community's voice will always be heard. We work with the community—we don't speak for them and certainly don't tell them what to do—and at the same time we provide a much-needed forum for end users, application developers, and system and distribution vendors to collaborate and continually enhance the Linux platform for their needs.

The Linux Foundation fosters the growth of Linux by complementing existing Linux achievements in these areas:

• Protecting Linux by providing legal services and sponsoring key Linux developers

It's vitally important that Linux creator Linus Torvalds and other key kernel developers remain



—continued on page 7

About the Author

Jim Zemlin, formerly executive director of the Free Standards Group, is the executive director of the Linux Foundation. He previously served as vice president of marketing for Covalent Technologies. Jim has also been a keynote speaker at industry and financial conferences and is an advisor on open source strategy to various companies and governmental groups including Hyperic, Zmanda and the Chinese Open Source Promotion Union.

ENTERPRISECLASS

Robust security. Fully scalable.
Java/J2EE environment.
Supports leading SQL databases.
In use today by Fortune 500.

ADVANCEDDESIGN

Best-of-breed CRM modules.
Runs on all standard-based platforms.
VoIP integration.
Customizable workflows.

OPENSOURCE

Source code freely available.
Flexible license.
Freely extend product capabilities.
Integrate with existing systems.

CUSTOMERMANAGEMENT

Manage sales pipeline end-to-end.
Leads management. Sophisticated
Help Desk. Custom surveys and
marketing communication.

No Matter the Weather, The Weather Channel® Is Open



The Weather Channel is open for business 24/7. When they decided to make their IT systems open as well, they turned to **Centric CRM**, the most advanced Open Source CRM system available. **Centric CRM** gave The Weather Channel's sales and support teams access to data from their patented technologies housed in 10,000 cable headends across the country. The result? Reduced total cost of ownership (TCO), improved customer view, and the elimination of multiple proprietary applications. Regardless of the conditions outside, **Centric CRM** helps The Weather Channel stay open.

Find out how Centric CRM can help your business get open:

Visit www.centriccrm.com or Call **757.627.8977**

Set up a free demo system / Become a Centric CRM Partner



Brian Shield
CTO, The Weather Channel

“Centric CRM is a key part of The Weather Channel's move to Open Source software, allowing us to mothball expensive, proprietary applications. Centric CRM is reliable, low cost and easy to use, adding feature rich functionality and flexibility to the way critical data is managed.”

Because Great Minds Are Open

 **CENTRICCRM**
Built for the Enterprise...Open to All

President & CEO

Fuat Kircaali fuat@sys-con.com

Group Publisher

Jeremy Geelan jeremy@sys-con.com

ADVERTISING

Senior VP, Sales & Marketing

Carmen Gonzalez carmen@sys-con.com

VP, Sales & Marketing

Miles Silverman miles@sys-con.com

Advertising Sales Director

Megan Mussa megan@sys-con.com

Associate Sales Manager

Corinna Melcon corinna@sys-con.com

EVENTS

Events Manager

Lauren Orsi lauren@sys-con.com

PRODUCTION

Art Director

Alex Botero alex@sys-con.com

Associate Art Directors

Abraham Addo abraham@sys-con.com

Louis F. Cuffari louis@sys-con.com

Tami Lima tami@sys-con.com

CUSTOMER RELATIONS

Circulation Service Coordinator

Edna Earle Russell edna@sys-con.com

SYS-CON.COM

Information Systems Consultant

Robert Diamond robert@sys-con.com

Web Designers

Stephen Kilmurray stephen@sys-con.com

Richard Walter richard@sys-con.com

ACCOUNTING

Financial Analyst

Joan LaRose joan@sys-con.com

Accounts Payable

Betty White betty@sys-con.com

SUBSCRIPTIONS

TOLL FREE 888-303-5282

201-802-3012

subscribe@sys-con.com

For subscriptions and requests for bulk orders,
please send your letters to Subscription Department

Cover Price: \$5.99/issue

Domestic: \$49.99/yr (12 issues)

Canada/Mexico: \$79.99/yr

all other countries \$99.99/yr

(U.S. Banks or Money Orders)

Back issues: \$12 U.S. \$15 all others

independent. The Linux Foundation sponsors them so they can work full time on improving Linux. The LF also manages the Linux trademark and offers developers legal intellectual property protection through such initiatives as the Open Source as Prior Art project, the Patent Commons (www.patent-commons.org), and sponsorship of the Linux Legal Defense Fund to deter and defend legal attacks on open source.

- **Standardizing Linux and improving it as a platform for software development**

A platform is only as strong as the applications that support it. The Linux Foundation will offer application developers standardization services and support that make Linux an attractive target for their development efforts. These include the Linux Standard Base and the Linux Developer Network. Currently all major distributions comply with the LSB and many major application vendors, like MySQL, RealNetworks and SAP, are certifying.

- **Providing a neutral forum for Collaboration and Promotion**


The Linux Foundation will serve as a neutral spokesperson to advance the interests of Linux and respond with authority to competitors' PR attacks. It hosts collaboration events among end users, application developers, the industry, and the Linux technical community to foster innovation and capture the viewpoints of its users. Through its workgroups, individuals can collaborate to solve pressing technical issues facing the Linux ecosystem in such areas as desktop interfaces, accessibility, printing, and application packaging.

I'd like to leave you with one important point. If you care about Linux, this is your foundation – it doesn't belong to me, or Linus, or big business, or the kernel developers. It belongs to everyone who cares enough to join and make Linux better. You can do this in many ways, not just by contributing or testing code. You can participate in one of our events and funnel your feedback to the community and to vendors who can make a difference. You can encourage your application developers to port their applications – if they haven't already – to Linux. You can tell your governments you want them to support true open standards in their software purchasing policy. Computing is entering a world dominated by two platforms: Linux and Windows. I've made my bet. The Foundation is here to make sure Linux unites its resources to challenge the privileged position proprietary platforms have enjoyed for too long. Please join us. 

from the editor —continued from page 3

what you are getting up front. Be wary of bait and switch tactics where you think you're getting a truly open source piece of software, when in fact you're getting a crippled or incomplete piece software. Some software looks to be open source but really is a gateway product to a commercially enhanced premium version. In my opinion, there's nothing wrong with commercial enhancements as long as there is value in the open source product sans any add-ons, and the line between commercial and community products are clearly drawn.

Why should you care about the spirit and intent of open source? I believe that for someone to really benefit from open source, they should be a participant in the open source community. If you aren't a good community member, it's unlikely that you will reap the benefits of the community, including word of mouth, collaboration, and other forms of user support. This is where I think the power lies to enforce or at least influence good open source citizenship. Public opinion and consumer pressure are effective methods for causing companies to behave in certain ways. I would urge users to consider this as they make their decisions.

How do you tell if a piece of software is really open source? First, check to see if the software is released under an OSI-approved license. Second, make sure that software that is open source offers some inherent value without third-party proprietary enhancements. Also look for vibrant forums and mailing lists with frequent posts and meaningful content. Usually the existing user base can clue you in to what you are in for. Finally make sure you are getting what you want. If you want open source freedom and value, examine your choices closely. 

Other Resources

- *The Open Source Definition:* <http://www.opensource.org/docs/definition.php>
- *The Open Source Initiative:* <http://www.opensource.org>
- *How Open Is "Open"?:* <http://opensource.sys-con.com/read/342346.htm>

This article is licensed under the Open Software License v 2.1:

<http://opensource.org/licenses/osl-2.1.php>.

BitLeap Offsite Backup

How open source software powers mission-critical backup services

by Guy Suter and Ian Berry

Learning how to solve existing problems is one of the keys to developing a successful product. Learning when to turn to the open source community to facilitate innovation is another quality that more and more companies are beginning to realize. The BitLeap story provides a real-world example of how new companies can innovate with open source and thereby create something new and useful.



About BitLeap

BitLeap is a provider of offsite backup services built on open source software. These services allow businesses of any size to implement automatic offsite data backups with ease.
www.bitleap.com

About the Authors

Ian Berry, executive engineer and co-founder of BitLeap, is the creator of the open source Cacti Project in use by thousands of organizations worldwide. He is a software engineer with experience in many Web technologies. Ian was instrumental in developing LeapServ, a Linux-based open source network backup appliance that remotely manages data backup and recovery.

Guy Suter, CEO and co-founder of BitLeap, has a mix of technical awareness and business administration experience. He has previously worked for large technology corporations like Qwest Communications and as part of executive management teams in regional technology companies. As CEO, he is responsible for ensuring BitLeap achieves complete customer satisfaction along with consistent innovation and growth in the data backup and storage industry.

As the founder of the popular open source Cacti project, BitLeap co-founder Ian Berry understood the merits of open source software. Hear how BitLeap incorporates open source to deliver cost-effective backup services.

First of all, interacting in the open source community comes with responsibility. Of course there's the responsibility of abiding by open source license stipulations and publishing code patches and alterations.

However, there's a more intrinsic responsibility – simply to share. Open source is about sharing code, of course, but also about sharing ideas, success, pride, and most importantly passion. This is our attempt to stick our heads out from behind our LCD panels for a moment and share all of those facets of our experience at BitLeap with a community we adore, and to demonstrate the power of open source in building the world's first remotely managed Linux-based network server, providing offsite data backup and recovery for SMBs and enterprise. At least we're pretty sure we were first.

We failed the Business 101 best practice advice of first identifying a problem your business will solve. In fact, we came up with the solution first and then started thinking about what problem we could solve using it. The exciting challenge for BitLeap is remotely managing servers spread across unpredictable networks thousands of miles apart under someone else's control at a low cost.

There is a long list of services that could be provided on this type of managed server platform. We choose data backup and recovery for business. This choice partly reflected our personal interest in tackling the challenge of developing an enterprise-grade data backup solution priced for SMBs. The choice was also motivated by the emerging market created by the salient rise in disaster recovery, offsite stor-

age, and the risks associated with traditional removable media for data backup.

Traditionally, data backups have been done to removable media devices like tapes, DVDs, and CDs. Linear backup approaches write down your entire data volume to one of these formats at specified intervals, usually once a day. Depending on the amount of data you're protecting, you might not be able to complete a full copy during a 24-hour period. Incremental backup techniques have helped get more granularity in the backup intervals but they have also increased the restoration hassle by requiring multiple tapes or devices to reconstruct the right data. The other problem is a lack of redundancy; if one tape or disk ever fails (and we know they can), you typically lose data. Advanced disk-based backup that included real-time offsite replication was available when we started out but the price was astronomical and the tools were mediocre.

We were inspired to start creating a new kind of data backup solution for businesses. This new solution would let IT administrators and consultants set up automated backups at frequent intervals, which would intelligently capture new and changed data on their network. It would provide an almost immediate local disk-based backup and securely replicate all of their data to multiple offsite locations on redundant RAID-5 storage servers by efficiently using the client's precious Internet bandwidth. Data restores would be easy and fast either directly from the local backup device or from the remote backup locations. Furthermore, we decided all of it should be done at prices that would attract SMBs as well as larger businesses.

Examples of Open Source Integration

Surprisingly the biggest challenge to providing offsite backup service is not the data storage requirements. While this is an important component, tracking the information about the data in a database quickly becomes the biggest challenge in scaling these kinds of solutions. To achieve our pricing goals for SMBs and have the ability to deliver a true enterprise-grade backup and recovery solution for a large business, we knew that larger RDBMSes such as Oracle and Microsoft SQL Server were simply not an option.

Gemini

2U

– Ultra Dense Series



Up to
**40% Less Power
Consumption**

Up to
40% Less Heat

- Two fully independent systems in a 2U
- Ability to run two discrete operating systems in one box
- Up to 8 CPU cores (forward compatible to 16)
- Up to 12 hot-swappable SATA, SCSI, or SAS hard drives
- RAID 0, 1, 5, 6, 10, 50 available on both systems
- Opteron™ or Xeon™ multi-core processors
- Up to 64GB memory per motherboard
- One available PCI-X or PCI-E slot per motherboard



Patent Pending

2 Systems in One

Built on open standards, the Gemini 2U elegantly accommodates two discrete motherboards in a 25" chassis uniquely designed for easy access from the rear.

Gemini 2U represents the realization of intoxicating power and superior environmental specifications, with up to 40% less power consumption, 40% less heat and 30% less noise. Remarkably, it all fits nicely into any standard rack. Front to back, Gemini 2U is both powerful and efficient.

At Open Source Systems we understand you need practical, customizable, and affordable solutions that are easy to manage and maintain.

For more information and to request your evaluation unit today, visit us at www.OpenSourceSystems.com, or call direct at 866.664.7867.



This is where open source RDBMSes such as MySQL and PostgreSQL come to the plate. Both of these products differ in their feature sets, performance, scalability, and licensing. As a result, BitLeap makes use of both powerful products to store and retrieve hundreds of millions of records while being able to scale to a continually growing (we hope) customer base. To make these database solutions truly scale, we write all of our code on top of a highly abstracted layer that handles database clustering, partitioning, and failover. We learned early on that you can only push a single clustered database system so far and that partitioning the data up into more manageable slices as it grows is the only way to scale while keeping costs down.

The challenge of remotely managing hordes of servers on distant networks without the luxury of being onsite, again led us back to open source. Without such a wide array of open source applications available, this task would be insurmountable. At the heart of this equation lies the extremely flexible and innovative Gentoo Linux distro. Gentoo's "do-it-yourself" approach provided an excellent framework to start with when rolling out BitLeap's remotely managed appliance servers (LeapSrvs) to customers.

All of these servers are birthed with PXE boot and are automatically built for each customer using a managed installer. Once built, these units automatically keep themselves up-to-date using an automated, almost organic distribution system run out of BitLeap's own little Capcom. All of the configuration information about every customer server is, of course, stored in the sky to enable easy recovery and replacement if needed.

The choice of programming language is especially important early in the development process since it dictates your target platform, scalability, and often the rate at which new features can be deployed. In the open source realm there are literally thousands of programming languages from which to choose. PHP came to our attention because of its increased mainstream acceptance over the past couple of years. We ultimately chose it because its straightforward syntax and flexibility allowed us to share our code base between our front-end Web applications and back-end applications.

It's fairly safe to say that BitLeap doesn't use PHP the way most companies typically do; that being primarily for driving dynamic Web sites. In fact, PHP can do so much more than simply outputting HTML markup. Using software written in PHP, LeapSrvs make network connections to Windows, Netware, Mac OS X, or Unix servers. The LeapSrv retrieves

the data to be backed up and breaks it down into small chunks for differential analysis. From there, the data is transferred offsite using a custom network protocol that's entirely encrypted, binary-safe, and capable of load balancing, rate limiting, and failover. On the server side of the equation, our PHP-based daemons pre-fork their children and handle hundreds of incoming connections per second from customer LeapSrvs. We do use PHP to output traditional HTML markup as well since our Web site and customer control panel share the same code base. This ensures tighter integration between the front-end and back-end and reduces duplication.

It's great to be able to rely on a high-level language like PHP for quick rollouts and tight collaboration. However, higher-level languages will never replace the performance optimizations achievable by using the lower languages like C or hand-tuning assembly code itself. Therefore, when the execution speed of a specific algorithm is critical to us, we rewrite the algorithm in optimized C with in-line assembly, and build it directly into the PHP source tree. This facility is simply not possible when using comparable closed source solutions

Real-World Impact

Perhaps the most exciting outcome of integrating open source software in new product development is seeing the end result of your new creation excel at solving real-world problems. At BitLeap, we have been delighted by positive feedback from our customers who often share how our data backup products and service save them time and money while providing more value than traditional removable media backup alternatives.

One such customer example is Motivait Seafoods out of Houma, Louisiana just south of New Orleans on the Gulf of Mexico. Motivait is a family-owned business run by people involved in the seafood industry since 1770. They use a revolutionary automatic shucking process that kills bacteria and at the same time creates the highest-quality oyster product available.

Motivatit owner Kevin Voisin says he realized it was time to revisit the company's backup strategy after witnessing the destruction caused by Hurricane Katrina. Mr. Voisin liked the ease of use and the fact that he had hardware he could carry with him away from a hurricane and simultaneously know it was being stored offsite. Motivait currently uses backup technology from BitLeap to protect data from its Windows servers.

It doesn't take a massive natural disaster to realize the benefits of having a solid, easy-to-use backup plan in place. Simple mistakes and

hardware failures are at the root of most actual data losses, as was in the case when Motivait recently lost its entire photo gallery data. This event gave Mr. Voisin a chance to see BitLeap's immediate restoration capabilities in action. According to him, "The recent restoration of our photo gallery happened so quickly it was hard to believe it had ever been lost."

It's not always practical for an organization to use open source software due to limited internal IT resources and the need for solutions with dedicated customer support. However, companies that are using open source integration to innovate new products with vast features for lower cost are providing a way that all businesses can benefit indirectly from the vast resources available through the open source community.

Giving Back

An important part of being involved in the open source community is giving back. There are lots of ways to do this. For instance, you can simply donate money to a project. It's important to remember that there are real people on the other side of every open source project out there who have put in a tremendous amount of effort. Most of them get little or no financial reward directly related to their efforts. Making a donation or buying something off of an Amazon wish list can go a long way towards motivating a developer to continue creating great software.

At BitLeap our need to make customizations to the PHP source tree or create open source code patches also provide us with a great opportunity to contribute back to the open source community. Submitting and publishing code modifications made to open source projects make the project itself stronger and can save time for other users with similar needs.

The best contribution to the development and adoption of open source software is to create and maintain new projects that solve existing problems. BitLeap's developers are encouraged to explore new ideas of authoring their own projects. As mentioned earlier, one of BitLeap's founding partners created the Cacti open source project, which has grown in popularity with people trying to simplify the management of network devices through centralized RRDtool-based graphing. We're now using the success of that project as a model to release some of BitLeap's internal tools used to manage our company under a new open source project. Maybe someday that will get even more attention than the commercial products we create, who knows? That's the beauty of open source in general. It's an ever-evolving paradigm where not even the original project authors know exactly where a project will end up or what impact it may have on the world. 

2007 VIRTUALIZATION CONFERENCE + EXPO

www.virtualizationconference.com



Register Now!
**EARLY-BIRD
SAVINGS!**
\$200

ROOSEVELT HOTEL
NEWYORK CITY

June 25-27, 2007

Virtualization: Solutions for the Enterprise.

Delivering The #1 i-Technology Educational and Networking Event of the Year!

As today's CIOs and IT managers rise to the challenge of using their enterprise computing infrastructure as cost-effectively as possible while remaining responsive in supporting new business initiatives and flexible in adapting to organizational changes, Virtualization has become more and more important.

A fundamental technological innovation that allows skilled IT managers to deploy creative solutions to such business challenges, Virtualization is a methodology whose time has come. The fast-emerging age of Grid Computing is enabling the virtualization of distributed computing, of IT resources such as storage, bandwidth, and CPU cycles.

But Virtualization can also apply to a range of system layers, including hardware-level virtualization, operating system level virtualization, and high-level language virtual machines.

Register Online!

www.VirtualizationConference.com



THE FIRST MAJOR VIRTUALIZATION EVENT IN THE WORLD!

**HURRY, FOR EXHIBITOR AND
SPONSORSHIP OPPORTUNITIES
CALL 201-802-3020**

*Learn from
the Experts...*

— BROUGHT TO YOU BY —



For more great events visit www.EVENTS.SYS-CON.com

COPYRIGHT ©2007 SYS-CON MEDIA. ALL RIGHTS RESERVED

ENTERPRISE
OpenSource
MAGAZINE

SOAWORLD
MAGAZINE

Next-Generation Linux Clusters

An innovative architecture for maximizing return on investment

by Bob Monkman

At a high level, a cluster is an aggregation of multiple standalone computers (nodes) linked together through a high-speed connection to create a single shared computing resource. A key benefit of this distributed processing architecture is that complex computations can be broken down and run serial or parallel among the individual nodes, resulting in a dramatic improvement in the time required to process large problems and complex tasks. For that reason, clusters are used for CPU-intensive jobs where massive processing power is required, such as when running simulations, computer-aided design, and rendering.

Clustered system architecture has steadily gained in popularity over the last decade for both high-performance and availability applications, partly because of the high return on investment (ROI) gained by using clusters. One reason for this high ROI lies with starting with dramatically lower capital costs. Most clusters are based on Linux, which provides a low-overhead licensing structure. Another reason many organizations choose clusters is the lower hardware costs of the most popular type of computational-intensive clusters, Beowulf clusters, given the economies of scales of x86-based servers. As will be seen as one reads this article, maximizing the ROI of Linux clustering is all about vastly simplifying the deployment and management of clusters to drive productivity up and operational expense down.

Thomas Sterling and Donald Becker developed the Beowulf architecture while at NASA. Beowulf clusters use standard commercial off-the-shelf (COTS) hardware, usually identical computers, running Linux. The nodes are networked into a small TCP/IP LAN and have libraries and programs installed that let processing be shared among them.

This type of clustered system architecture combines the power of many COTS systems to form a single more powerful system and can reach and exceed the performance of traditional SMP and vector-based supercomputers. But it's also far less expensive to reach that same performance level. For example, \$1 million can buy well over 1,000 processors using leading-edge AMD- or Intel-based serv-

ers. But you can't even approach that "bang for a buck" with shared-memory multiprocessor (SMP) systems. The cost savings means individual departments can afford their own "supercomputers." By some estimates, overall costs for a COTS-based cluster may be a third to a tenth the price of a traditional supercomputer.

Using COTS technology also helps users take advantage of several trends, including improved technology (like microprocessor performance and memory density) and commodity networking, especially gigabit Ethernet, soon to be followed by 10Gb, (which makes it possible to design distributed-memory systems with tolerable bandwidths and latencies). The result is a compelling order-of-magnitude cost savings for a very powerful and far more accessible (for the typical user) high-performance computing (HPC) center.

Next-Generation Clusters

Many scientific and technical applications moved to first-generation clusters in the late 1990s, including fluid dynamics, stress and deformation analysis, seismic and reservoir modeling, large-scale operational research, molecular and protein modeling, and financial analytics. More recently, the digital content creation market has become primarily Linux cluster-based. Before clusters were widely used, these applications ran on vector computers, on large-scale SMPs, and, to a lesser extent, on massively parallel systems (MPPs). These classes of machines shared high initial investment costs and training of personnel and maintenance often doubled or tripled the initial capital expenditure. Clusters significantly reduced capital costs as well as ongoing administration costs.

There was another more substantial difference between Linux clusters and other options available at the time, though. The first Linux clusters were designed so that end users could gain full control of their tools and have access to supercomputer performance without the costs or burden of dealing with such complex and expensive platforms. Biochemists, fluid dynamics experts, mechanical engineers, and others in research labs and

About the Author

Bob Monkman is director of software product management, Penguin Computing.

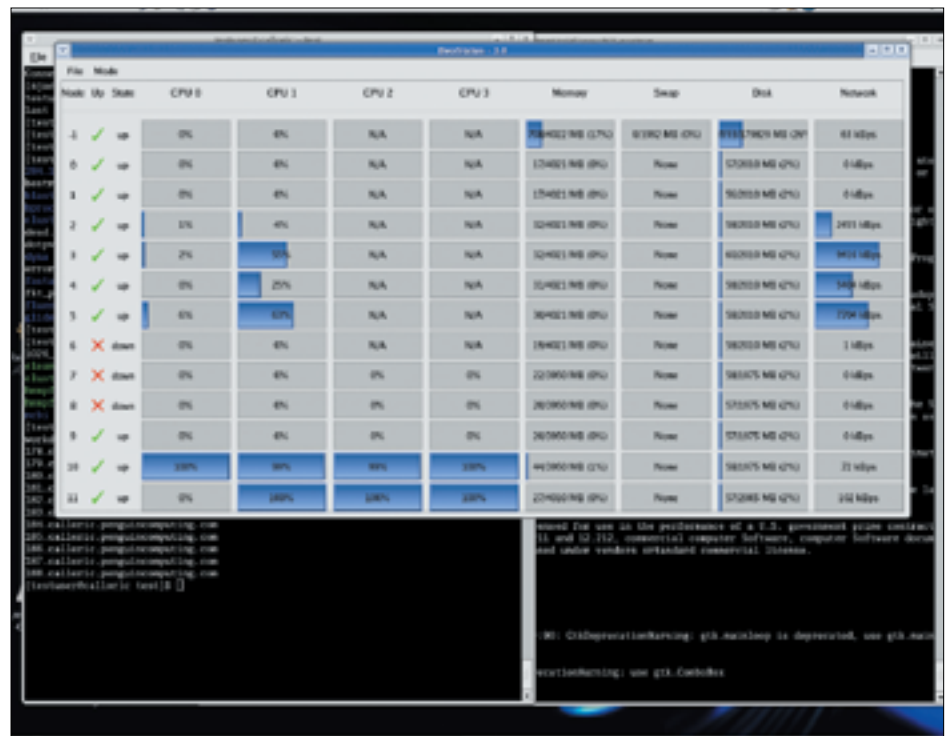


Figure 1

universities were at the forefront of this silent revolution — not computer scientists. Linux clusters were being designed by the users for the user.

By the end of the '90s, however, Linux clusters started to become interesting to Fortune 500 companies as well, expanding beyond the relatively small arena of research and education. However, the first-generation cluster architecture that came out of the Beowulf project at NASA only addressed the problem of getting machines to work together, making sure the networking code worked, and building the communication library layer. There were several areas where first-generation clusters needed improvement, not the least of which was ease of use for the less technical user.

One key improvement next-generation clusters had over first-generation designs was the illusion of a single system that makes management easier. In a necessarily simplified fashion, some key elements that were required to set up an easy-to-use and administer cluster by creating the illusion of a single system include:

- A stable, reliable, and observable booting system
- Mechanisms for initial provisioning and updates of all system elements
- A “stateless” system software provisioning model in which the operating environment underlying the applications loads directly to memory

- A consistency model for handling dynamic updates
- A lightweight compute operating environment that's stripped of unnecessary software that would otherwise consume resources, introduce latencies, and open security holes
- A single point of monitoring and control for both systems and jobs
- A file system or other storage management that suits the application

Some non-commercial cluster systems create a partial single system illusion by requiring network virtual memory or a consistent global file system, or implementing transparent process migration. However, these designs handle failure poorly because the system must go through a time-consuming lock recovery process or even kill all processes related to the failed machine if any of the nodes fail.

Because end users built the first generation of clusters, they continued to build or grow their own ad hoc clusters as the needs of their organizations dictated. Ad hoc do-it-yourself clusters are cumbersome to use and generally adapted only to batch processing. But usage habits change. HPC clusters have been effectively adapted for batch computing environments but, until recently, have not been well suited for interactive use. This is why batch schedulers have been so prevalent — the issue for users is that they are presented in a

traditional cluster with 100 separate servers for a 100-node cluster. Managing jobs and their related processes, data, and other things means manually interacting with all of those machines, a tedious and time-consuming process.

This complexity is even worse for the admin who must set up and maintain the cluster. The admin is presented with 100 separate machines that must be loaded with the operating system and cluster management software and this can take five-30 minutes for each machine if everything goes well. Then all of the servers must be configured for access, users, security, and coordinating as a cluster. Extensive scripting is usually required and different people must maintain these scripts over time adding overhead and cost.

What's more many do-it-yourself solutions are simply open source projects and lack important formal support and documentation. There are many associated costs in terms of time and resources that project-level solutions carry over the lifetime of the cluster. In addition, most clusters get different parts of the cluster stack from several different suppliers so customers face the additional challenge of who to turn to for support when problems occur.

It's far more cost-effective to use a commercially developed and fully supported solution that can simplify the management of large pools of servers, leveraging the best of open source and value-add proprietary software, and wrap it all together in a well-architected, pre-integrated, documented, and supported solution. And that's what happened after the initial development of clusters. Next-generation clusters architected from the ground up to be "commercial-grade" don't face any of these issues listed above, benefiting both commercial and non-commercial users.

In addition, as the market matured and more complete solutions have come available, the value proposition for Linux clusters has expanded to include significant total cost of ownership (TCO) savings spanning hardware, software, and support costs; flexibility in configuration and upgrade options; freedom from constraints of single-vendor development schedules and support; greater flexibility through open source software customization; and rapid performance advances in industry-standard processor, storage, and networking technologies.

Maximizing Return on Investment

Next-generation clusters, almost all of which are commercially developed, offer several specific areas of additional value compared to previous generations. This is due in part to optimization of the initial first-generation Beowulf architecture into software that's

properly architected, very tightly integrated to work out-of-the-box, and offers additional features that vastly simplify the deployment and management of large pools of servers. By examining the benefits of some of the architectural concepts of next-generation systems for compute clustering in more detail, we can understand their advantages.

Faster, Better, Cheaper Deployment

It's dramatically easier and faster to provision and manage large server pools when you don't install a full operating environment on the hard disks, the path chosen by some leading-edge cluster software vendors. A full OS installation to the disk drive is relatively slow, generally taking 15 to 30 minutes to complete depending on a variety of factors. This can be scripted or even parallelized but it'll still take hours to provision an entire cluster once all the preparations are completed.

In some next-generation systems, though, the installation is completely integrated and needs to be done only once and only on the designated Master node, regardless of the size of the cluster. The compute nodes are then auto-provisioned with a cluster-aware pre-configured operating system environment directly to memory and no further configuration is generally required since the cluster software sets up the entire configuration. This process takes approximately 20 seconds for each node and they are ready to run.

The value of rapid stateless provisioning is even more apparent during ongoing cluster operation when compute nodes fail, are updated, or are re-provisioned or new nodes are added. The stateless approach is effortless and nimble in comparison to a traditional local full-install. A full disk-based install on a replacement or new server is going to take the same 15-30 minutes as an original install and will generally have to be scheduled by one's IT staff, when they can get to it. This increases the time required to restore the cluster to full operation.

In a next-generation architecture, software updates only need be applied to the Master, which will then auto-update the compute nodes quickly and on-demand. The method used in most architectures, however, requires elaborate scripting at best and increases the chance that something can go wrong part-way through. This will result in another problem, version skew. The correctness of the applications that run on multiple compute nodes is often dependent on everything being precisely the same on each processing element. The tiniest difference in a driver or library can render the results useless after days of calculation — if the application runs at all.

“Businesses that ignore the potential of SOA will find themselves outpaced by rivals who improve their agility and transform themselves into new kinds of enterprises

— Yafim Natis, Gartner Analyst

3-DAY EVENT!

SOAWorld

Plus **2007**

Enterprise OpenSource

Conference & Expo 2007

TOPICS INCLUDE:

SOA Web Services

- > AJAX and SOA
- > Web 2.0
- > Universal SOA
- > Protecting Web Services
- > Troubleshooting SOA
- > Governance
- > Open-Source SOA
- > XBRL
- > Service Virtualization

Open Source

- > Open Source Business Models
- > Open Source ESB
- > OpenAjax Alliance
- > SaaS and Open Source
- > Spring, Hibernate and Eclipse
- > Seam
- > Open Source Penetration
- > Monetizing Open Source
- > Open Source Databases
- > AMQP
- > Open Source Middleware

June 25-27, 2007
Roosevelt Hotel / New York City

Register Online! www.SOAWorld2007.com

11th International
SOAWorldTM
CONFERENCE & EXPO

2007 is to many industry insiders shaping up to be a major inflection point in software development and deployment, with SOA, Web Services, Open Source, and AJAX all converging as cross-platform and cross-browser apps become the rule rather than the exception.

Accordingly the 11th International SOA Web Services Edge 2007 again seeks to offer comprehensive coverage and actionable insights to the developers, architects, IT managers, CXOs, analysts, VCs, and journalists who'll be assembling as delegates and VIP guests in The Roosevelt Hotel in downtown Manhattan, June 25-27, 2007

Co-located with the 2nd Annual Enterprise Open Source Conference & Expo, the event will deliver the #1 i-technology educational and networking opportunity of the year. These two conference programs between them will present a comprehensive view of all the development and management aspects of integrating a SOA strategy and an Open Source philosophy into your enterprise. Our organizing principle is that delegates will go away from the intense two-day program replete with why-to and how-to knowledge delivered first-hand by industry experts.

Visit soaeosconference.sys-con.com for the most up-to-the-minute information including...
Keynotes, Sessions, Speakers, Sponsors, Exhibitors, Schedule, etc.

2nd Annual
ENTERPRISE > 2007
OPENSOURCE
CONFERENCE+EXPO

SOAEOSCONFERENCE.SYS-CON.COM

REGISTER ONLINE TODAY

SAVE \$200!

(HURRY FOR EARLY-BIRD DISCOUNT)

BROUGHT TO YOU BY:



» **SOA World Magazine**
focuses on the business and technology of Service-Oriented Architectures and Web Services. It targets enterprise application development and management, in all its aspects.



» **Enterprise Open Source Magazine**
EOS is the world's leading publication showcasing every aspect of profitable Open Source solutions in business and consumer contexts.

SYS-CON EVENTS For more great events visit www.EVENTS.SYS-CON.com

Exhibit and Sponsorship Info:

Call 201-802-3020 or email events@sys-con.com

Depending on local full-install operating environments is one sure way to create problems of version skew. In the end, you may end up spending hours reinstalling everything fresh and getting back to a known state. Local copies get stale and inconsistent causing wasted time and costly rework. However, the latest generation of cluster software generally provisions to the nodes exactly what is running on the master and inherently manages dependencies and versioning so there's guaranteed consistency.

Greater Performance

One of the immediate benefits of the lightweight compute nodes mentioned as key components of next-generation clusters is performance. Part of the reason that modern compute nodes can be provisioned in less than 20 seconds is that the OS is significantly smaller. Since compute cluster nodes aren't general-purpose machines, they don't need most of the software provided in a full Linux distribution.

Related to performance is improved memory use, especially important if insufficient memory forces an application to swap space. On a typical cluster compute node, next-generation clusters can begin with a very small footprint of memory (16MB) and dynamically add support only as required as compared with 400MB for a full static Linux installation.

A more significant issue is the scheduling latencies that many of the standard Linux services can introduce over long-running applications. It's been shown that these scheduling latencies can cut the cluster performance of real-world applications by 5%-50% as cluster configurations scale out and they can be impractical to isolate since they are very application-dependent.

Enhanced Scalability

Scale-out performance is greatly improved due to the unique design of these new architectures. Some designs employ a single primary daemon on compute slaves and leverage this daemon to run jobs, get standard I/O, and logs and statistics out on those slaves, all from the Master node. Enhanced scalability occurs because compute nodes can be added on-demand and common tools can be optimized to leverage the support already built into the architecture. As long as the networking and storage infrastructure is designed correctly, there's no bottleneck created in the system to compromise scalability.

Simpler, Faster Administration by Design

Another way in which modern clusters enable scalability is tied to their ease of use. By using virtualization, they make large pools of servers act and feel as if they were a single consistent virtual system. At least one next-


generation cluster software package creates "single system image" behavior with the Linux you already know. It does this by extending the Linux configuration on the Master node to have a single unified process space. From both the administrator and the user's point-of-view, a 100-node cluster with 400 processors appears very much like a 400-processor machine.

So the compute servers are fully transparent and directly accessible if needs be. However, if you're interested in the compute capacity presented at the single Master node, you need look no further than this one machine. The simplest way to convey how this might manifest itself to the typical user is to offer the example of the everyday task of issuing the ubiquitous "ps" process list command. After issuing the command, what you get back is a listing of all processes running on all machines as if they were just one machine. You can still tell which processes are running where in the cluster with a simple addition to the command, but only if you really care to know it. Other standard Linux commands work in the same intuitive way as on a single machine.

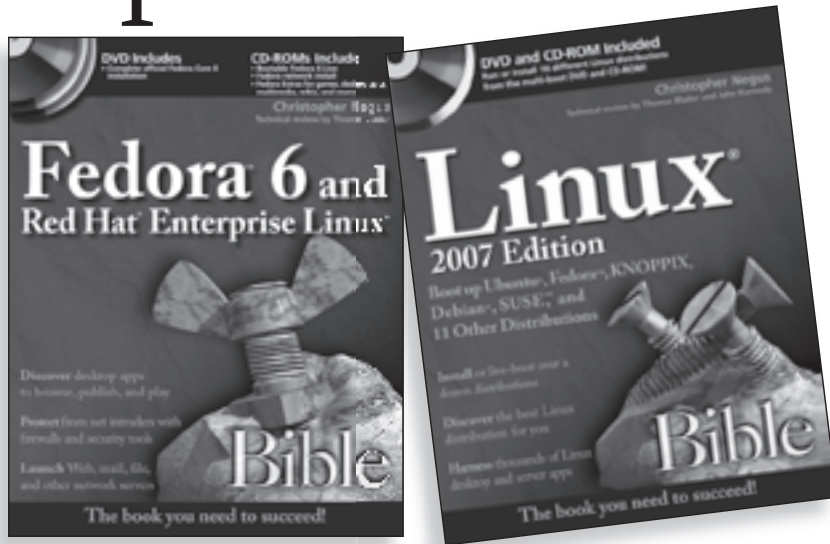
When you want to add a user and set up passwords, you only have to do it on the Master. When you want to run a job, you run it on the Master and simply tell it how many processors you need (even non-MPI jobs). If you need to terminate a job, it's done on the Master node and automatically removed cleanly on the compute node(s) it actually resides on. Of course, you can run jobs or general commands on specific nodes if you have to. If you need to see the vital statistics of load, memory usage, disk usage, etc. on any or all nodes, add one command line or GUI invocation on the Master node and you get it.

Conclusion

The latest generation of commercial clusters re-architects the foundation of cluster architecture using several well-recognized concepts in a unique combination that delivers virtualized cluster systems that make large pools of servers appear and act like a single consistent virtual machine. A properly architected solution that leverages stateless provisioning and a lightweight compute operating environment that simulates the appearance and end-user experience of a single virtual system has a tremendous ripple effect on rapid provisioning, manageability, scalability, security, and reliability within the cluster.

The result is an elegantly simple and powerful new paradigm of virtualized clustered computing. This new paradigm eliminates the need for multiple levels of cost and support. It also dramatically increases efficiency and reduces operating costs while delivering a dependable HPC service to organizations in highly competitive business environments. 

When Chris Negus speaks, people learn Linux!



978-0-470-08278-2

978-0-470-08279-9

Get the latest Linux Bibles by Christopher Negus.
The books you need to succeed.



Wiley and the Wiley logo are registered trademarks of John Wiley & Sons, Inc. Linux is a registered trademark of Linus Torvalds.

Looking to Stay Ahead of the *i*-Technology Curve?

Subscribe to these
FREE Newsletters >

Get the latest information on the most innovative products, new releases, interviews, industry developments, and *i*-technology news

Targeted to meet your professional needs, each newsletter is informative, insightful, and to the point.
And best of all – they're FREE!

Your subscription is just a mouse-click away at www.sys-con.com

IT solutions

NET JOURNAL

JDJ

WebServices JOURNAL

information STORAGE+ SECURITY JOURNAL

LinuxWorld JOURNAL

wireless JOURNAL

LINUX BUSINESS WEEK

XML JOURNAL

MX developer's journal

WebSphere JOURNAL

wldj JOURNAL

COLOFUSION JOURNAL

SYS-CON MEDIA

The World's Leading *i*-Technology Publisher

Trust But Verify

Eliminating insider threats as an essential security and compliance best-practices measure

by Ellen Libenson

Some years ago, a small manufacturing firm was hijacked by its IT employees. They didn't use weapons, but they did commit murder: They killed the business. Unfortunately, they were aided and abetted by the company's own lax security policies.



The firm had ignored the best practice of segregating duties: The IT employees who developed the core manufacturing software were also tasked with administering the system, supporting the end users and managing the backup process. When business conditions required the company to reorganize its IT staff, some of the employees who were to be laid off took revenge by sabotaging the mission-critical application. They erased audit trails, reassigned user rights, and even destroyed the backup tapes. The manufacturer never recovered.

While this may be an extreme example of the havoc supposedly “trusted” insiders can wreak, it's not an isolated instance. Most privileged users have earned their organizations' trust, but a single rogue employee can do great damage, disrupting or destroying systems, or stealing corporate data. Even if these acts don't leave a company in ruins, they can expose a business to heavy financial losses, fines for compliance violations, and negative publicity.

Many organizations have yet to pay close attention to combating this threat. Companies in recent years have focused on the external threats from hackers to phishers that the media has underscored. Yet privileged users in organizations can represent a more significant risk. Their credentials often supply them with the “keys to the kingdom” that enable them to access and view data that they shouldn't, and access and manage IT systems in a way that can damage operations. Sometimes, the damage is an accident, but it's equally likely to be malicious.

Some startling facts about the risks posed by insiders are revealed in “The Insider Threat Study: Computer System Sabotage in Critical

Infrastructure Sectors.” The May 2005 report, published by the U.S. Secret Service National Threat Assessment Center and the Carnegie Mellon Software Engineering Institute's CERT program, focused on nearly 50 cases of insider sabotage carried out between 1996 and 2002. It found that 86% of the insiders who engaged in sabotage were employed in technical positions, most as system administrators. Most had privileged access to systems when hired, though less than half had authorized access at the time of the incident. Fifty-seven percent of the insiders exploited systemic vulnerabilities in applications, processes, or procedures, or some combination of them. Eighty-one percent of the organizations affected suffered financial losses — in at least one case tens of millions of dollars. The business operations of three-quarters of the affected organizations suffered a negative impact, ranging from deleted sales records to destruction of proprietary software, and the reputations of nearly 30% were damaged.

Beyond all this, organizations expose themselves to penalties if they fail to ensure the integrity of their financial systems and statements and protect customer data as required by regulatory mandate, including Sarbox, HIPAA, and state privacy laws. The 2006 CSI-FBI security study reveals that respondents consider data protection, regulatory compliance, and identity theft as their most critical computer security issues over the next two years. But failure to follow best security practices, including limiting system access by privileged users, compromises corporate compliance efforts. Failing to do so may even personally affect the organizations' key executives, who must vouch for management's

About the Author

Ellen Libenson, vice president of product management at Symark Software (www.symark.com), a provider of security administration solutions for heterogeneous IT environments. elibenson@symark.com

having adequate internal controls over financial reporting.

There is certainly reason to worry about compliance measures being compromised on Unix and Linux systems, where the root user account gives system administrators access to all the power of the root. As more and more organizations elect to run mission-critical financial, CRM, and other applications on heterogeneous servers, controlling privileged accounts is more crucial than ever.

Without proper controls, anyone with access to the root account — the virtual “keys the kingdom” — is given complete super-user privileges without justification based on their job classification, duties, or role in the IT department. This violates security best practices by exposing proprietary systems and information to malicious activity and sabotage that could result in catastrophic information leakage or mistakes that could bring down a network.

Action Must Follow Awareness

The good news is that there's a growing awareness of the threat posed by privileged users, particularly at financial institutions, which have seen some of the more high-profile breaches.

Indeed, nearly 70% of respondents to the CSI-FBI security study attribute some percentage of corporate losses to insider threats. Seven percent believe that insiders account for more than 80% of their organization's losses.

Further proof of awareness comes from the 2006 E-Crime Watch survey, conducted by CSO Magazine, the U.S. Secret Service, the CERT Coordination Center, and Microsoft. Respondents to that survey who experienced e-crime said that, to the best of their knowledge, 63% of intellectual property thefts involved insiders, as did 56% of other proprietary information thefts, 49% of sabotage cases, and 71% of incidents having to do with the intentional exposure of private or sensitive information. Fifty-five percent of the organizations that have experienced security events now report at least one insider event, up from 39% in 2005.

Indeed, vendors that provide identity and access management products for Unix and Linux security and administration, including Symark Software, have seen sales rise in direct response to company needs to comply with privacy laws and federal regulations.

Now for the proverbial bad news: even as companies are becoming more aware of the problem, many still aren't taking some of the most basic steps to lock down their environments. On Linux and Unix systems, big strides can be made simply by putting technology in place that delegates the root's abilities to system

administrators without providing the root password and providing indelible audit trails of user actions.

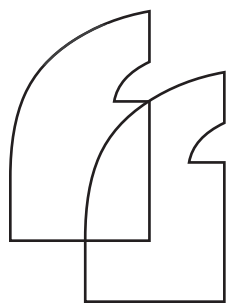
So what exactly holds companies back from such a step? Misreading the potential cost of the risks they're running, for one thing. For instance, some healthcare organizations may consider that the \$250,000 fine for a HIPAA privacy violation doesn't justify the price of purchasing and deploying another security solution. They miss the bigger picture costs here, by not considering all the other expenses they will incur should they suffer a client data breach. That includes notifying their customers that their data files have been breached, a cost of about \$15 a pop. Those fees can add up quickly.

The “it can't happen here” factor also plays a role. There's a real human desire to want to trust the people we work with. Unfortunately, that noble instinct runs up against the darker side of human nature — people being tempted to take revenge for perceived slights, bad performance reviews, or layoffs; to steal and sell company data to pay off gambling debts; to misappropriate company funds to cover personal expenses.

Finally, organizations may face some internal pushback on security efforts since they disrupt the IT status quo. Many IT professionals see themselves as the organization's heroes in white hats, coming to the rescue when viruses are on the loose or systems go down. They may view plans to limit their root user accounts as an affront to their personal ethics, as well as an obstacle to being able to do their jobs effectively. They may also believe that audit trails are meant to spy on whether their time is well spent. The situation may be complicated by politics — the demand for implementing security and administration tools may be driven by security officers or internal auditing teams rather than IT operations leaders.

The situation can be prickly, given the tensions that sometimes exist among these groups, and the trend towards downsizing IT staff. You can hardly blame system administrators for fearing that their authority is being curtailed, or that their productivity is being monitored. Yet, at the same time, IT employees understand better than most the temptations to explore unauthorized data. According to a recent survey by the analysts at Dark Reading, approximately 10% of IT/security pros say they abuse security privileges to access unauthorized data on a regular basis. Nearly 30% are not regular offenders, but they admit to having accessed unauthorized data during their careers.

Since making any changes to infrastructure or operations goes more smoothly when IT leaders get buy-in from staff, it's important to address their concerns. For one thing, system adminis-



Unix and Linux systems, where the root user account gives system administrators access to all the power of the root, are a security hazard

trators should be made aware that the steps are being taken not just to ensure the proper use of privileged access in the department, but so that access privileges can't be abused by personnel outside the IT department. The more people who know the password to the root account, the harder it is to manage the password, and the more opportunities there are for non-IT personnel to find it out. Joe in accounting may see the password left on a Post-It note on a system administrator's desk, for example, or Jane in facilities might hear it used in a conversation over a speakerphone. No one could argue with plans that will put an end to those threats.

It might also be pointed out that audit trails of system administrators' actions are equally valuable to prove how diligently a worker is performing for the company.

The ultimate argument, of course, is that compliance and legal requirements demand that the business take steps to limit Unix and Linux system administrators' access to the full power of the root. As much as anything, such a deployment makes a case for job security, since a serious breach of systems or data that results in heavy financial losses, fines, or legal expenses could mean lost jobs.

The Fix Is In

Segregating duties at a process level is a key step that organizations must take to limit risks from insider threats that may compromise compliance requirements or otherwise damage the business. Technology gives businesses the teeth they need to ensure that the process is respected. Removing access to the root password and allowing root powers to be shared selectively, specifying exactly which program each user can execute as root, is critical. That will make it impossible for the system administrator charged with mounting disks and modifying tapes to take over the duties assigned to other administrators, such as modifying corporate databases or adding users. And, since the root password is never given out, the risk of it becoming freely available, even to non-IT users, is eliminated.

Equally important is to be able to trace administrative actions, but not simply by turning on the Unix accounting feature. That function

records every activity on the system — not just root actions — thus creating an overload of raw data (and IT groups are already dealing with enough of that). Even worse, traces of actions are susceptible to being erased from the account logs by any user who has root access. An indelible audit trail, maintained on a separate machine from the one the system administrator has access to, is key to knowing who took what administrative actions, when, and on which machines.

Implementing an indelible trail that records who initiated the action, on what date and time, and on which machine and directory also serves as a deterrent. If privileged employees know their activities are readily traceable, they will be less likely to attempt to engage in malicious activity. Having a complete record of a session is equally valuable for discovering innocent mistakes that result in system problems and help organizations avoid those issues in the future.

As important as it is to be able to record keystrokes to replay an entire root session and see exactly what a user typed, it's equally critical to be able to designate forbidden keystrokes that might, for example, bring down a server. Such a feature protects against even the highest-level privileged users at an organization from acting on rogue intentions, since entering those keystrokes will require approval from multiple parties.

Businesses should also consider the risk that rogue employees will attempt to access root privileges after-hours. On its own, Unix doesn't make it convenient to restrict access to the root to certain times of the day or certain days of the week. Frankly, the number of methods a determined rogue user may use to gain access to the Unix root password or exploit it are shockingly high. For instance, he may monitor unencrypted user input and output control messages sent over the network or create backdoors and log in from remote locations with a stolen password.

Some organizations may think they are protected against damage with an open source product that provides some root delegation capabilities. But these tools still enable anyone with any system administrator responsibilities to access the root password, and they lack indel-

ible audit trails and keystroke logging. Those capabilities are critical requirements for really addressing the threat posed by disgruntled insiders, particularly in Unix, which has historically lacked emphasis on secure system administration.

Sudo or Su-Don't?

After deciding to implement controls to protect the root account, organizations face the challenge of how exactly to do it. Because most IT administrators are trustworthy, organizations must properly manage access rights without hampering the IT staff. One of the fundamental questions that must be thoroughly addressed is whether to deploy commercial solutions or rely on freeware.

Perhaps the most widely used freeware for delegating responsibilities in Linux and Unix environments is Sudo. The basic intention of Sudo is to provide administrators with a way to let users access certain programs that require the root password without giving them complete root privileges. While Sudo contains a handful of positive attributes, IT managers familiar with the program understand that its drawbacks make it an insuffi-

cient solution.

In a small environment — for example, 10 or 20 servers — Sudo may be able to handle a business's access control needs. However, larger organizations often have hundreds of servers running dozens of different versions of Linux and Unix operating systems. Sudo fails catastrophically when trying to scale to an infrastructure of this magnitude.

With Sudo continuing to fall out of favor with administrators at large-scale IT environments, organizations are turning to commercial identity and access management solutions as more effective way of addressing insider threats to satisfy compliance regulations and follow best security practices without alienating the IT department.


Be On the Lookout


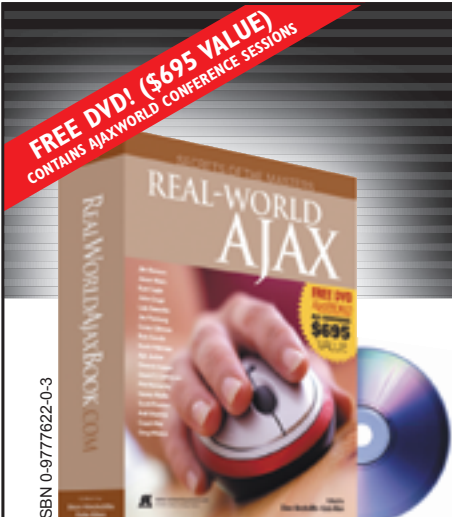
The bottom line: It's okay to trust your employees — so long as that trust can be verified. Organizations that put solid technology and policies in place that limit privileged users' access to the systems and data for which they are responsible, and record indelible audit trails of their actions, make it extremely difficult for internal staff to run amok.

No technology, of course, will ever substi-

tute for people's ability to pick up signs that there may be a problem with an employee—one that could lead him to commit sabotage, steal data, or otherwise harm the organization. The Insider Threat Study, in fact, found that 80% of the insiders exhibited unusual behavior in the workplace prior to carrying out their malicious activities.

At that now-defunct manufacturing company, one of the IT employees involved in the sabotage had long been expressing his unhappiness about the impending layoffs and making veiled comments about his intentions. But they were ignored. The lesson to learn is this: As important as it is to deploy technology to stop threats in their tracks, it is equally important for managers to seriously consider subtle (or not so subtle) threats an employee may be making, and for HR departments to provide support services for employees who are in financial trouble or are facing a job loss.

It is better to forestall trouble than to ask for it. Taking such steps, and deploying technology solutions that pre-empt malicious activity, will help businesses ensure that no rogue actions will cause harm to their operations. 





JOIN THE AJAX REVOLUTION!

Secrets of the Masters: Real-World AJAX

Edited by Dion Hinchcliffe & Kate Allen

"If you're looking for a one-stop shop for an AJAX book...you would have a long search to find a better overall resource than what you find in the chapters of this book."

Order Online at RealWorldAJAXBook.com and get **40% OFF** Regular Bookstore Price!

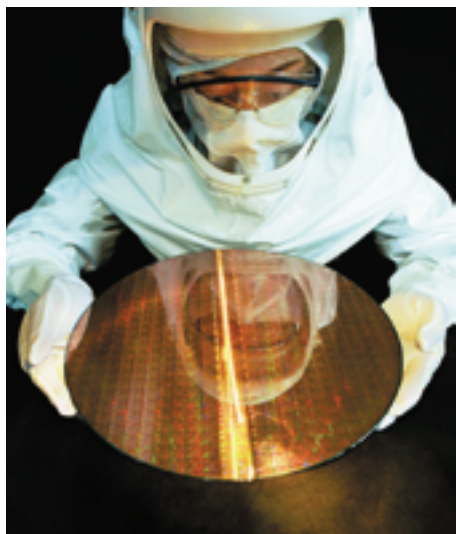
 books.sys-con.com
from the World's Leading i-Technology Publisher © COPYRIGHT 2007 SYS-CON MEDIA

Multi-Core Debugging and Performance Enhancement

Additional pressures on complex applications

by Douglas Eadline and Vince Hauber

Computer systems have fully entered the age of multi-core processing. This trend was examined in a white paper entitled, "Preparing for the Revolution, Maximizing Dual-Core Technology." While much of the focus had been on dual-core processors, Intel has now delivered quad-core processors and AMD has announced quad-core availability in 2007. Software developers who are just taking advantage of multiple processors also need to be looking ahead since even larger multi-core platforms will reach the market soon. This new technology puts additional pressures on complex applications. In essence, the push beyond dual-cores has made the choice of development tools a key factor for successful projects.



About the Authors

Douglas Eadline has over 25 years experience in high performance computing. Further information about multi-core and high performance computing can be found at <http://basement-supercomputing.com>.
deadline@Basement-Supercomputing.com

Vince Hauber, a senior product manager with Concurrent Computer Corporation, has over 40 years experience in system software and platform solutions. Concurrent is a leading provider of real-time Linux distributions and tools.
vince.hauber@ccur.com

When writing multi-core software, there are two important issues facing software developers. The first and most familiar is correctness of the program, i.e., does the program work properly and get the right answer. While this seems obvious, in multi-core applications, incorrect program behavior can arise from several sources not present in single-core design.

The second important feature is program performance. An assumption that speed-up is automatic just because a program has been modified to utilize multi-cores can be a huge mistake. Improved performance is not necessarily automatic or guaranteed.

With the introduction of quad-core processors, applications have the option of using eight cores on one dual-socket motherboard. Proper use of these additional resources will be critical to software development. This article will provide some insight into these issues and present some solutions to the problems facing multi-core programmers using the Linux operating system.

Prerequisites

As suggested in the white paper mentioned above, some consideration should be given to the method of parallelization. Two basic models are available to programmers – program

threads and message passing. While message passing is the predominant method used in High Performance Computing (HPC), threads are becoming more important as both the core and socket sizes increase in most new motherboards.

Given that message passing programming is covered by other sources, this article will focus on threaded programming on multi-core processors using Linux.

A general approach to parallel computing is first to take a working sequential program, a program that is known to operate correctly on one CPU, and use it as the basis for a parallel program. This technique has two big advantages. First, the program can be tested to ensure that it operates at an acceptable level of correctness. (i.e., the program works and any issues are known in advance). Second, the programmer has a performance baseline with which to measure the new multi-core program.

If a new program is being written, it's highly advisable to create and debug a sequential version first. Adding the additional complexity of parallel programming to standard software development makes isolating bugs and problems more difficult.

In general, threaded programs often have multiple independent parts that share the same data set. One way to use threaded programming is to allow subroutines to operate independently and at the same time. It's also possible to break large loops up into threads as well. In all cases, however, a good design will allow the threaded version to "collapse" into a single thread by setting a thread count variable to one.

Finally, since there are adequate resources for learning about threaded programming, we'll concentrate on ways to ensure correct program operation and optimal performance.

Programming with Threads

The thread model is a way for a program to split itself into two or more concurrent tasks. These tasks can be run on a single processor in a time-shared mode, or on separate processors

(e.g., the two cores on a dual-core processor can each run threads). The term thread comes from “thread of execution” and is a reference to how a fabric (computer program) can be pulled apart into threads (concurrent parts). Threads are different from individual processes (or independent programs), because they inherit much of their state information and memory from the parent process. Threads provide the ability to share memory and offer very fine-grained synchronization with other sibling threads.

On Linux and Unix systems threads are often implemented using a POSIX Thread Library (pthreads). There are several other thread models (Windows threads) from which the programmer can choose, however, using a standards-based implementation like POSIX is highly recommended. As a low-level library, pthreads can be easily included in almost all programming applications. This article will discuss debugging and optimizing Linux threads.

Ensuring Program Correctness

Of course, all programmers know that just because a program finishes, it doesn't mean the answers are correct. With multi-core programming, this situation may manifest itself differently. The same program and data, run at different times, may produce different answers, or cease functioning altogether for no apparent reason. We'll explore some of these causes and present some solutions.

Deadlock and Livelock

Because individual threads share memory and I/O, there must be a way for each thread to protect memory segments and I/O resources while they're using it. To solve this problem, threaded code employs “locks” that provide the thread exclusive access to a resource.

While the introduction of locks provides a solution to this problem, it also creates the possibility of a deadlock occurring in the program. Consider the following scenario where there are two global variables called Total and Sub-Total.

Process 1 locks the variable Total

Process 2 locks the variable Sub-Total

Process 1 attempts to acquire a lock on Sub-Total

Process 2 attempts to acquire a lock on Total

The program is now in a deadlock situation. Deadlock may develop every time the program runs or it may happen randomly when the timing is right. There's also a condition called “livelock” that's similar to deadlock. Livelock results when a thread or threads are constantly trying to acquire a lock, but can't because it's being used by another part of the program (or more likely hasn't been unlocked properly). Livelock may stop certain threads, but not necessarily cause the

whole program to stop. With livelock, a program may appear to be working, but in reality it's often stuck in a loop, which can make it hard to determine the exact cause of the problem.

While deadlock and livelock can often be recognized through program behavior (i.e., the program stops, crashes, or gets stuck in one place), locating the cause can be difficult at times unless you can look at the threads in real-time. Certain lock conditions have a temporal component and can be hard to reproduce unless the program is running in production mode (i.e., a specific sequence or timing of events caused the situation).

Identifying and fixing deadlock and livelock situations requires a good debugger. In today's multi-core multiprocessor environments, efficient debugging requires the user to navigate between processes and threads from a common environment without disturbing the program running. A debugger should operate with minimal application intrusion so that deadlock and/or livelock can be identified and fixed in real-time (e.g., applications hot-patched while the application is running). Users should also be able to stop a thread and observe the state of all the other threads running in the application.

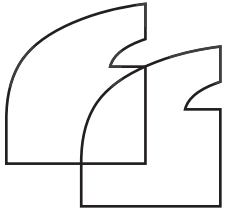
In some cases, the events leading up to a lock situation may not be immediately apparent to a debugger or the exact timing will have to be recreated without a debugger present. To resolve these issues, it's often necessary to generate program execution traces so that actual events and interactions between threads can be examined.

A graphical event analyzer tool is ideally suited to this purpose. With such a tool, user programs and the Linux kernel can be traced live or viewed in a post-execution file. In both cases, user and kernel events can be collected from the multiple processes executing simultaneously on multiple cores and processors, and then examined to determine what caused the lock.

Race Conditions

A race condition is similar to a deadlock but can cause more subtle problems. Suppose the programmer in the example above only allows Process 1 to own the lock on Total to avoid deadlock. Now suppose all program threads can lock and add to the Sub-Total variable. If Process 1 decides to add the Sub-Total to Total at some point, it may not have all the contributions to Sub-Total. These types of situations can be hard to identify and result in the program providing different answers when using the same data.

A race condition is particularly difficult to find. Often they result from subtle timing assumptions in the program that aren't guaranteed to be true each and every time. The program will occasionally provide the wrong result. Moreover,



The push beyond dual-cores has made the choice of development tools a **key factor for successful projects**

when the programmer adds some debugging code or attaches a simple debugger and steps through code, the problem goes away. This type of behavior is often called “non-deterministic” since it “just happens” from time to time. To find these kinds of problems, a debugger must be able to examine program execution in real-time.

Again, a robust Linux debugger that allows real-time modification and monitoring of an executing program is essential to resolving race conditions. There's no need to add code or alter program behavior to find and repair problems. Another useful debugging feature is the ability to analyze and display memory allocations and de-allocations. A race condition, and deadlock/livelock for that matter, can occur when a thread behaves unexpectedly and memory management issues are often the source of such problems.

Mismatched Communication/Synchronization

As simple as it sounds, mismatched communication is often the cause of many parallel computing problems (both with threads and message passing). Communication is an essential part of parallel programs. (If parts of the program don't communicate then it would not be a parallel program!) Communication can also be used as a synchronization method. For instance, a thread may wait until another thread tells it to start or stop a task. This kind of communication is often called “blocking” and provides the programmer with some synchronization points.

The other method of communicating is called “non-blocking.” In this method a task may be busy computing and occasionally checks to see if a message has arrived. While this asynchronous behavior can be more efficient than a synchronous blocking approach, it's also subject to communication deadlock and race conditions.

In either case, a communication mismatch occurs when a sender or receiver isn't available. These situations can occur by outright programming error or by using asynchronous methods like “non-blocking” communication. Similar to race conditions, asynchronous communication issues can be non-deterministic.

One way to debug communication/synchronization issues is to use a multi-thread event analyzer. This tool lets system and user-requested events be logged from multiple processes executing simultaneously on multiple processors or cores. The result is a minimal-overhead, high-resolution trace of your application behavior.

Taking Debugging to the Next Level

All of these conditions require the programmer to debug the programs deeper than ever before and at the same time touch the program very lightly. Many standard debuggers aren't prepared for this level of interaction. Indeed, the problems associated with multi-core debugging usually require attacking the problem using both debugging and tracing methods. Desirable debugging features are the ability to debug multiple processes in a single session, allowing runtime program modification and patching. Beyond interactive debugging, both data and event tracing offer the lightest way to touch a multi-threaded program.

Program Performance

There are many parameters that contribute to overall program performance. On multi-core systems, it's important to sort out the performance issues by application and by thread. In addition, multi-core versions of your software should work faster than a single-core version. Often programmers are stumped when they achieve poor performance on parallel systems. Having the right tools to drill down into the processor and cores is essential if bottlenecks are to be identified.

Load Balance

The most obvious issue with multi-core processors is to make sure all your cores are busy and work is balanced across them. An unbalanced application will result in poor performance improvement because one core (or thread) has become the slow step in the process. For this reason, it's important to know “what's running where?” and “how resources are being used” to eliminate bottlenecks and ensure optimum performance.

Besides load balance, there are other parameters that can indicate how well a core is performing. Monitoring imbalances in parameters like context switches, interrupts, memory paging, and processor affinity can also indicate poor performance.

A Linux performance tuning utility can aid significantly in system and application tuning. It's ideally suited for multi-core application analysis because it allows the user to probe by process and thread while observing key system metrics that can influence performance. Information can be displayed for individual processor cores so that the real-time load balance can be observed while the application is running.

Viewing Performance

A program running on single-core processors can be viewed as a progression of events dictated by the programmer. Fleshing out performance issues often requires programmers to look at the events in detail. Multi-threaded codes running on multi-core environments often have deeper issues that could not exist on single-core processors. These deep issues often involve timing between threads and the resources they touch. One way to go deep into program behavior is to instrument a program so that detailed performance data can be harvested. In preserving exact runtime behavior, user-written instrumentation can be difficult and introduce further issues in the code (or hide issues as well). It's often best to use a tracing library that's been specifically designed for this purpose.

During program operation, an instrumented program should emit trace data that's collected with minimal influence on the running program. After the program finishes, a trace file that contains the runtime information will provide insights into program and thread behavior. Another important area that is often neglected in program tracing is the operating system. Often serious bottlenecks can be traced to certain tunable aspects of the Linux kernel and thus are hidden from the end user.

Trace files can produce extremely large amounts of information when multiple cores are involved. So is it important to use a tool that can sort through the trace information easily so that critical points can be identified.

As previously mentioned, a Linux trace utility that's aware of both user application and kernel-level activity can provide a level of information not obtainable with ordinary debuggers. This information can also be used to enhance the performance of the application. Users generally know the critical parts of their programs and the trace tool can provide a detailed window into these areas. A well designed GUI front-end can also provide easy navigation through the trace data.

Recommendations

When developing multi-core code, the following recommendations will assist in faster and better code generation.


- 1) Create a reference sequential program that can be used as a baseline for both results and performance.
- 2) Use your reference program as a basis for your threaded program and make the number of threads tunable so that your program can be easily collapsed to one thread.
- 3) Use a real-time debugger that preserves the timing of your program. Additional tools be-




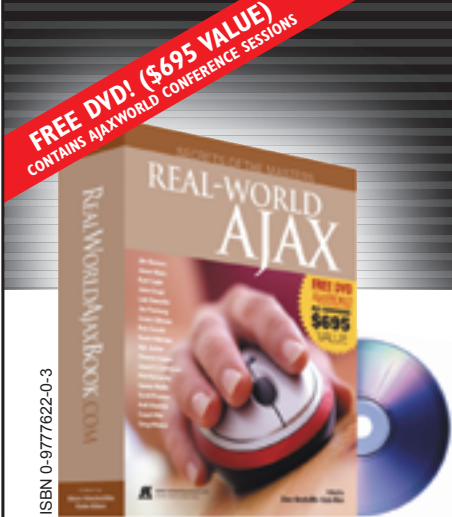
yond debuggers that provide data and event tracing may be needed to ferret out difficult issues.

- 4) Once the application is running correctly, examine the load and resource use by using a system/application tuner and event analyzer.
- 5) As a final step, check your assumptions. The current x86 64-bit architectures all run the same binary codes, but have widely varied hardware implementations. Some time invested analyzing your program with regard to these areas will be well spent.

Conclusion

The speed increase offered by multi-core designs has become an exciting part of software creation. When writing and debugging code for multi-core systems, however, attention must be paid to new issues that weren't present before. Lock and race conditions can be particularly difficult to resolve with tools designed for single processors. Furthermore, improved performance isn't always guaranteed and may require deeper analysis of runtime behavior than was needed in the past.

Finally, it can't be stressed enough that as hard as single-processor programming can be, multi-core can be much more difficult. A good set of tools is essential to controlling costs and achieving delivery schedules. 




JOIN THE AJAX REVOLUTION!

Secrets of the Masters: Real-World AJAX

Edited by Dion Hinchcliffe & Kate Allen

"If you're looking for a one-stop shop for an AJAX book...you would have a long search to find a better overall resource than what you find in the chapters of this book."

Order Online at RealWorldAJAXBook.com and get **40% OFF** Regular Bookstore Price!

 **SYS-CON BOOKS** books.sys-con.com
from the World's Leading e-Technology Publisher © COPYRIGHT 2007 SYS-CON MEDIA

How Open Is 'Open'?

Industry luminaries join the debate

by Jeremy Geelan



In order for company to describe itself as an “open source” company, need it merely help you make the switch to open source in your company – or does it have to be one that lets users freely download, compile, and use the software in question? Where is the dividing line?

How open is “open”?

Nat Torkington raised the very real question of whether the term “open source” is now completely meaningless, i.e., its meaning has now been sucked out of it by companies that purport to be open source yet don’t allow users to freely download, compile, and use the software in question. One example Torkington cites is SugarCRM, whose license he describes as “a questionably modified OSI-approved license.”

But Sam Minee of SilverStripe didn’t agree with Torkington and argued that, in order to describe itself as an “open source” company, all it need be is “a company that will help you make the switch to open source in your company.”

Where is the dividing line? How “open” is open? And is “open source” as a term dead or dying? At *Enterprise Open Source Magazine* we contacted a range of FOSS luminaries for their take on the issue.

“I’D SUPPORT THE REUNIFICATION OF THE TERMS ‘FREE’ AND ‘OPEN SOURCE.’”



“I see open source as a term relevant to the way communities function and I’d support the reunification of the terms ‘free’ and ‘open source’ around the concept of free software being developed in open source communities. On that basis it’s not dead.”

– **Simon Phipps**, Chief Open Source Officer (Sun Microsystems)

“THE OSI HAS DONE A GREAT JOB PROTECTING THE DEFINITION OF OPEN SOURCE”



“Open source is freely downloadable, usable, and redistributable by its community of users and contributors. While trademarks and certification value-add should be and is protected, open source has no oddball gated community, look-but-don’t-touch, private source, or badgeware restrictions. There are no open source, on-ramp ‘childrens’ editions’ to closed source products offered by a true open source company. People and companies should get credit for their contributions and work in an open manner in the community, but not with restrictions on the openness of the code. To date, the OSI has done a great job protecting the definition of open source. Let’s hope they continue that tradition.”

– **Pierre Fricke**, Director, Product Line Management (SOA Products, Red Hat)

“THERE ARE LEVELS OF OPENNESS...”



“The OSI’s Open Source Definition, first drafted by Bruce Perens for Debian in June 1997, is the most obvious place to turn to when attempting to apply an ‘open source litmus test.’ However, using the OSI’s criteria, a number of efforts we generally consider ‘open source’ would fail this test. There are levels of openness and I don’t believe a strict interpretation is going to be practical. It’s up to each person to determine what ‘open source,’ or more broadly, ‘open’ means within its specific context. The challenge right now is that it’s hip to be ‘open,’ even if you’re not. Are there some ‘posers’? It all depends on your ‘open’ threshold.”

– **Raven Zachary**, Analyst (*The 451 Group*)

“UNDERSTANDING OPEN SOURCE IS AN EVOLUTIONARY PROCESS”



“Throughout the time I’ve been involved with open source, I have been fascinated by the recurring pattern of companies wishing to align themselves with open source principals while they simultaneously maintain their essentially proprietary business instincts and models. At OSI we have seen that the process of growth in open source is more evolutionary than revolutionary. We invite public debate with each successive wave of newcomers to start the process to close the gap between what they imagine open source to be and the reality of what is required (and why).”

– **Danese Cooper**, Board Member (*Open Source Initiative*)

“THE CONCEPT OF OPEN SOURCE IS FAR FROM DYING”



“The whole point of open source licensing is that no one company can do anything to threaten the freeness and openness of the software. The open in ‘open source’ derives its meaning from the idea, embodied by the GNU General Public License, that the freedom to receive, distribute, alter, and copy software is guaranteed. As long as people understand that software licensed with such mandates constitute open source, and that all other licenses don’t, there is no concern about the meaning of ‘open’ shifting. The concept of open source is far from dying. The software that has been released as open source (or free software as some prefer) must be distributed with the very licenses that make them free. Until the world has no use for Linux, Drupal, PostgreSQL or the Apache server, open source is alive and well.”

– **Robert Douglass**, (*the Drupal Association*)

“IT’S PERFECTLY FINE IF COMPANIES DO NOT WISH TO MEET THE OPEN SOURCE DEFINITION”



“Open source is open source, and it was defined by the OSI for a reason – to prevent brand dilution of the term open source. It’s perfectly fine if companies do not wish to meet the open source definition, as they are certainly free to define another term that better reflects their strategy. It’s better for the software industry on the whole if the taxonomy waters are not muddied more than they already are.”

– **John Mark Walker**, Community Manager (*Hyperic*)

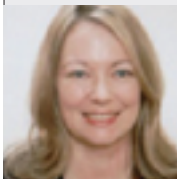
“NAT TORKINGTON IS RIGHT”



“I very much welcome any company that will ‘help make the switch to open source,’ and I wouldn’t want to rule any one of them out by stating that they are incorrect in using the term in their marketing campaigns. But Nat Torkington is right that the term ‘open source’ is becoming more and more meaningless as it is no longer just a plain and simple definition of how a piece of software is developed and maintained. For this reason, there has been the suggestion within the Apache Software Foundation to use the more specific term ‘open development.’ By referring to ‘open development,’ a company (or rather, community) very clearly states that it is developing and releasing its software in a very specific way that allows anyone to download, build, and run the code. And that is exactly what the more abstract wording of ‘open source’ is currently unable to do.”

– **Arjé Cahn**, CTO (Hippo)

“I WOULD LIKE TO SEE MORE CONSOLIDATION OF LICENSES”



“Open source isn’t black and white. Individuals choose to consume open source for different reasons. Companies and individuals decide to invest their time and money in open source development for different reasons. The proliferation of new open source licenses does cause unnecessary confusion and complexity for both individual consumers and vendors building solutions. I would like to see more consolidation of licenses to make it easier for users to understand. We selected OSI-approved licenses for our open source initiatives to make it easier for users to adopt and modify the software to meet their needs.”

– **Debbie Moynihan**, Director of Open Source Programs (IONA Technologies)

“THE OPEN SOURCE DEVELOPMENT MODEL IS A MERITOCRACY”



“I think the question ‘How Open Is Open?’ is perhaps the wrong question as it ultimately will answer itself. One of the advantages of the open source development model is that it’s a meritocracy. So by virtue of the quality of your software and the effectiveness of your open source strategy, the ones that do things right will be successful. It doesn’t matter what you or I or other community leaders think, it’s the opinions of the end user that matter. To that end, companies that cloud the open source issue up front with a pseudo open source strategy will fall by the wayside, and those that are true innovators will flourish. I believe this whole thread was sparked by a current trend of ‘open source crippleware’ where ISVs put out a minimalist version and then hold back the useful additions to create a business that coerces users to a paid subscription. In the short term, it has worked for a number of open source vendors; in the long term, I think these strategies might be less successful. In my own experience I have been intimately involved with two open source projects heavily over the last year with both types of licenses. The first open source project I was involved with significantly was NetDirector, which is released under the MPL with an attribution clause, and the second is Zenoss, which is licensed under the GPL. The one that has been more successful in terms of adoption is GPL-licensed Zenoss; that’s shaped my thinking and my recommendation to others.”

– **Mark R. Hinkle**, VP of Community and Business Development (Zenoss Inc.)
and Editor-in-Chief, Enterprise Open Source Magazine

"THE DEFINING CHARACTERISTICS OF 'OPEN'..."



"The defining characteristics of 'open' are:

- A genuine, recognized open source license. The Apache License, which Spring uses, is very liberal and does not lock out any commercial use
- Visible repositories: free information about what is developed and by whom
- Open issue tracker: free information about bugs and resolution, and the ability for anyone to file issue reports
- A real community, perhaps funded by one or more companies but with wide participation
- The ability for anyone to gain committer status through merit and dedication"

— **Rod Johnson**, CEO and founder (Interface21)

"COMPANIES THAT DO NOT PROVIDE 'OPEN SOURCE' SOFTWARE SHOULD NOT BE CALLED OPEN SOURCE COMPANIES"



"Even though I use the term Free Software more often to emphasize the freedom aspects, I think the term open source also has a strong meaning from a commercial point of view. The open source definition defined by the OSI clearly defines whether a license can be considered open source 'compliant' or not. However, we must make sure that no company or individual defines a license as open source when it is not. Companies that do not provide 'open source' software should not be called open source companies. If they are producing a mix of proprietary software and open source software, then I don't see a problem with it, even though they don't totally respect our freedom."

— **Damien Sandras**, Creator and Developer
(Ekiga's VoIP and videoconferencing software)

"OPEN SOURCE BY ITSELF DOES NOT DESCRIBE TODAY'S SPECTRUM OF OPEN SOURCE PROJECTS"



"Commercial open source product companies are continuing to edge away from the 'official' definition of open source, perhaps because it doesn't provide a way to prevent a new competitor from using the original source code as the basis for a new, competitive offering. In addition, while most of the commercial open source companies follow an open source distribution model, very few of them follow an open source development model that includes community contributions. So I agree with Nat that the term 'open source' by itself does not describe today's spectrum of open source projects. I routinely distinguish 'commercial open source' from 'community-based open source,' where the former includes those companies trying to generate revenue from open source software and the latter are non-commercial, including 'free/libre' software."

— **Tony Wasserman**, Executive Director,
Center for Open Source Investigation (Carnegie Mellon West)

THREE REASONS TO

blog-n-play.com

1 Get instantly published to 2 million+ readers per month!

blog-n-play™ is the only **FREE** custom blog address you can own which comes instantly with an access to the entire i-technology community readership. Have your blog read alongside with the world's leading authorities, makers and shakers of the industry, including well-known and highly respected i-technology writers and editors.

2 Own a most prestigious blog address!

blog-n-play™ gives you the most prestigious blog address. There is no other blog community in the world who offers such a targeted address, which comes with an instant targeted readership.

3 Best blog engine in the world...

blog-n-play™ is powered by **Blog-City™**, the most feature rich and bleeding-edge blog engine in the world, designed by Alan Williamson, the legendary editor of **JDJ**. Alan kept the i-technology community bloggers' demanding needs in mind and integrated your blog page to your favorite magazine's Web site.



www.TAMI.linuxworld.com

"Many blogs to choose from"

PICK YOUR MOST PRESTIGIOUS ADDRESS

IT Solutions Guide	MX Dev. Journal
Storage+Security Journal	ColdFusion Dev. Journal
JDJ: Java	XML-Journal
Web Services Journal	Wireless Business &Tech.
.NET Dev. Journal	WebSphere Journal
LinuxWorld Magazine	WLDJ: WebLogic
LinuxBusinessWeek	PowerBuilder Dev. Journal
Eclipse Dev. Journal	

3 MINUTE SETUP

Sign up for your FREE blog Today!





SWsoft to Package Virtualization with SUSE Linux Enterprise

(Hannover, Germany) – SWsoft has announced it will deliver its Virtuozzo server virtualization software bundled with SUSE Linux Enterprise Server 10 from Novell.

SWsoft will take a standard distribution of SUSE Linux Enterprise Server and add the Virtuozzo software to deliver one integrated solution to customers. As a result, customers will have another choice for virtualization that leverages SUSE Linux Enterprise 10, an interoperable open platform for mission-critical computing.

SWsoft plans to make the Virtuozzo and SUSE Linux Enterprise Server 10 bundle available worldwide in the second quarter. It will run on servers using X86, AMD64, EM64T, and Itanium processors. Support for the software bundle will be provided by SWsoft and backed up, as necessary, by Novell for all the unmodified components of SUSE Linux Enterprise Server.

With SWsoft Virtuozzo, a single Linux operating system instance can be dynamically partitioned into multiple, highly efficient, and scalable virtual environments. Virtuozzo includes patented technology that enables density of up to hundreds of virtual environments on a single physical server.

Unlike virtual machines and hypervisors, Virtuozzo addresses the challenges faced by today's data centers of operating system sprawl. Its unique architecture and management tools make it the ideal solution for provisioning, monitoring, and managing virtualized server resources.

www.swsoft.com



Xandros Linux Server to Provide Enterprise-Grade O3Spaces OpenDocument Collaboration

(Hannover, Germany) – Xandros, Inc., a provider of easy-to-use Linux desktop, server and management solutions, and O3Spaces, B.V., provider of O3Spaces Workplace collaboration for OpenOffice.org, StarOffice and MS Office, have announced an agreement to provide OpenDocument and MS Office document collaboration, management and retention services for the forthcoming release of Xandros Server 2.0 – Standard Edition. O3Spaces Workplace provides any office user with a professional out-of-the-box extension for team and project collaboration, and meets the growing demands from governments and companies for an affordable, platform independent and integrated collaboration alternative to the Microsoft Office SharePoint server.

www.xandros.com www.o3spaces.com



Centric CRM Releases Open Source Centric CRM 4.1

(Norfolk, VA) – Centric CRM, developer of an open source customer relationship management (CRM) system, has released Centric CRM 4.1, featuring a number of improvements that add both power and productivity to enterprise CRM. Centric CRM 4.1 introduces Action Plans, a new tool that allows a company to establish workflow best practices, as well as expand support for a total of 11 leading database systems.

Other new features and upgrades in Centric CRM 4.1 include:

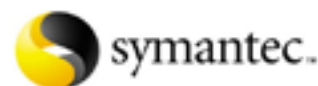
- Support for both Geronimo and IBM WebSphere Application Server-Community Edition J2EE application servers
- A new Web Content Management System (beta) that users can leverage to build and manage a fully functional external Web site that is seamlessly integrated into their CRM system
- Web 2.0 collaborative capabilities such as discussion forums and project management
- Asterisk for VoIP voice capability, enabling appropriate customer contact personnel to be notified by screen alert whenever an account calls in as well as outbound calling to account contacts at the click of an icon
- Portal access to account information for external customers, allowing customers to access their account details, trouble tickets, and project plans via the Web.
- Certification on the entire line of IBM hardware and processors up to and including mainframes.

www.centriccrm.com

Symantec Enables Data Center Standardization with Support for Red Hat Enterprise Linux 5

(Cupertino, CA) – Symantec Corp. has announced its portfolio support for Red Hat Enterprise Linux 5, as well as new storage management capabilities for Xen virtualization. Through its partnership with the Linux market leader, Symantec plans to provide its core Veritas Storage Foundation, Veritas Cluster Server, and Veritas NetBackup solutions for customers as they adopt the new subscription release of Red Hat Enterprise Linux.

Symantec and Red Hat are also announcing a collaborative effort to deliver highly trained support expertise to customers through joint engineering, pre-release testing, and a unique and innovative shared virtual lab support model.



To further aid customers in managing their complex data centers, Symantec is leveraging the open source Xen server virtualization that will be bundled as part of Red Hat Enterprise Linux 5 and other Linux distributions to launch a beta program for its Storage Foundation and Dynamic Multi-pathing technologies. For the first time with any server virtualization platform, this will unite best-in-class storage virtualization and I/O multi-pathing with virtualized workloads, enabling a standardized, centralized management model across all workloads and all operating systems, physical or virtual.

Symantec and Red Hat are unveiling a joint customer initiative, which will ensure faster, more comprehensive customer support resolution to better serve enterprise organizations. The innovative cooperative agreement between Symantec and Red Hat provides technical support engineers remote access to Symantec and Red Hat testing lab environments, combining both companies' server and storage infrastructures to address virtually all customer support requirements. In addition, technical support engineers are being cross-trained to better support joint customers. The result will be superior support quality, integrated escalation processes, and much faster time to resolution since the partners can troubleshoot and fix customer issues together on the same hardware.

www.symantec.com

Concurrent Delivers High-Performance Linux Solution to Toshiba TEC

(Duluth, GA / Tokyo) – Concurrent, a provider of time-critical Linux operating systems and integrated computer solutions for mission-critical applications, has announced that Toshiba TEC has selected Concurrent simulation technology for the development of its next-generation document processing systems. Toshiba TEC is a Japanese-based manufacturer of document processing, telecommunications, and retail information systems.



Toshiba TEC chose a Concurrent iHawk solution running Concurrent's RedHawk Real-Time Linux to assist in the development of new Multi-Function Peripherals (MFP), devices that perform a variety of functions such as copy, print, fax, and scan. Toshiba purchased multiple iHawk high-performance PCI-based computer systems featuring AMD Opteron 64-bit processors. Their next-generation MFPs are developed using numerical model simulation. Distinct particle modeling will be used to analyze toner levels.

The Concurrent RedHawk solution simulates the behavior of more than 10,000 small toner and ferrite particles in a magnetic field. The required dynamic calculation interval is less than 100 nanoseconds and calculated for one complete revolution of a magnetic roller. The system helps determine design parameters such as particle size and distribution, roller speed, magnetic field strength with minimal experimental (laboratory) work.

www.ccur.com

Emulex Open Source Driver Now Shipping with Red Hat Enterprise Linux 5

(Costa Mesa, CA) – Emulex Corporation has announced the availability of its Linux driver for Emulex LightPulse Host Bus Adapters (HBAs) as part of the new Red Hat Enterprise Linux 5 distribution. As a result, Emulex LightPulse HBA customers will be able to immediately benefit from the new virtualization, clustering, and manageability enhancements available in the Red Hat Enterprise Linux 5 distribution.



Key features and support from Emulex for Red Hat Enterprise Linux 5 include virtualization support ensuring interoperability with the Xen 3.0 server virtualization module; clustering support with multipathing compatibility enabled by the Red Hat Enterprise Linux 5 Device Mapper module; increased manageability by fully implementing "sysfs" permitting third-party applications to collect HBA subsystem information.

The Emulex Linux driver is fully qualified and supported by major server and storage OEMs including HP, IBM, EMC, Dell, and HDS. For enterprise customers, this means comprehensive support from both Red Hat and their storage solution provider. The Emulex Linux driver works across all models of Emulex LightPulse HBAs. The Emulex HBAnyware management application supports Red Hat Enterprise Linux 5 and customers can leverage it to simplify HBA management for the entire SAN, saving time and money. Emulex is working closely with a number of Linux and Xen technology providers, offering a development kit that enables the creation and management of multiple logical connections servicing virtual servers or individual applications.

www.emulex.com



Novell Delivers Automated Desktop Management Solution with Native Directory Support for Microsoft Networks

(Hannover, Germany) – Novell has introduced a new solution to ease desktop management of Windows operating systems, including Vista, for customers who need streamlined management in their Microsoft and Novell environments. With native integration for both Microsoft* Active Directory* and Novell eDirectory, Novell ZENworks Configuration Management enables real-time, identity-based systems management that speeds service delivery and ensures accurate desktop configuration. This approach makes desktop configuration changes immediately whenever network user accounts are changed. Built with the Novell Desktop-to-Data Center Management Blueprint for delivering ITIL-based services, ZENworks Configuration Management improves IT efficiency and increases user productivity.

ZENworks Configuration Management gives organizations the ability to manage desktops based on their user's network identities, which can include a combination of user business role, location or group membership. This enables the delivery of applications, content and IT resources tailored to the personal needs of each user. For customers who require device-focused management, the product can also function independently of eDirectory and Active Directory.

www.novell.com



Open-Xchange Unveils Community Project – To Build Next-Generation Open Source Collaboration Suite

(Tarrytown, NY) – Open-Xchange Inc., a provider of open source collaboration software, has announced the opening of the Open-Xchange Community project and the availability of source code and development documentation for its latest AJAX-based e-mail and groupware technology.

Open-Xchange and the Web hosting company, 1&1 Internet, recently announced the availability of MailXchange, 1&1's hosted e-mail and collaboration solution based upon Open-Xchange Inc.'s Smart Collaboration technology. Open-Xchange has released the source code of 1&1 MailXchange server components under General Public License (GPL) and the AJAX-based user interface (UI) under Creative Commons License "Attribution-NonCommercial-ShareAlike 2.5".

www.open-xchange.org



HP Reduces Customer Administrative Costs and Wait Time with Blade Network Virtualization Technologies

(Palo Alto, CA) On the heels of taking the No. 1 position in factory revenues and units shipped for blade servers,(1) HP today introduced virtualization technologies that simplify the connectivity and management of its HP BladeSystem c-Class architecture, the company's core solution for the next-generation data center.

HP also announced the launch of its first workstation blade for the c-Class blade portfolio, the HP ProLiant xw460c Blade Workstation, to provide users with increased performance and scalability.

Dramatically simplifying network connectivity and common server management tasks across the data center, the HP Virtual Connect Ethernet and Fibre Channel virtualization modules enable administrators for the first time to wire just once and then add, replace or recover servers on the fly. In addition to lowering capital and operating costs, the modules deliver simplified management capabilities that reduce the time it takes to execute these actions from hours or days to a matter of minutes.

"HP is creating systems that dramatically reduce installation time, cut costs for administration and reduce barriers to change within an IT organization," said Ann Livermore, executive vice president, Technology Solutions Group, HP. "With HP Virtual Connect, CIOs and IT administrators are set to overcome two of the biggest problems of today's data center: administrative productivity and the need to drive rapid change."

The new technologies bolster the HP Adaptive Infrastructure offering, which helps customers move towards 24x7, "lights out" data centers, and will add to the company's momentum in the blade server market. According to analyst firm IDC, HP surpassed IBM to take the lead in factory revenues and units shipped for blade servers in the fourth calendar quarter of 2006.(1)

"Customers are on a treadmill of complexity and cost that limits the value IT is able to deliver to the business," said Zane Adam, director, Windows Server Division, Microsoft Corp. "HP's new Virtual Connect modules with HP BladeSystem c-Class, running Windows Server 2003 Enterprise Edition, help customers operate mainstream business services more nimbly and cost-effectively, increase the manageability of their IT environments and improve responsiveness to new business requirements."

hp.com/go/setITfree40



LINUXWORLD[®]

CONFERENCE & EXPO

it360[°]
CONFERENCE AND EXPO

IT360° is a powerful multi-sector experience.

Get the IT360° advantage – a single source of knowledge, realistic strategies, and tools to help you solve the critical burning issues today, paving the road for efficiency and productivity tomorrow. www.it360.ca.

Conference: April 30 – May 2, 2007

Trade Show: May 1 – May 2, 2007

Metro Toronto Convention Centre – TORONTO, CANADA

IT360° Conference and Expo is an ITWorld Expo event produced by ITWorld Canada the trusted name in Information Technology Resource Media.

www.it360.ca



Next-Generation AMD Opteron Processors Power New Line of Open Source Systems Server and Storage Solutions

(Sunnyvale, CA) Open Source Systems, formerly Open Source Storage, a leading provider of enterprise server and storage solutions for mission critical business data centers, announced today it will incorporate Next-Generation AMD Opteron™ processors into its new server and storage product line. With today's announcement, OSS customers will continue to benefit from AMD's recognized leadership in x86 dual-core processing performance and energy cost-savings for multi-processor capable servers and workstations.

"Our new line of double-density Gemini servers will usher in a new era of ultra-dense, power efficient, data center workhorses." Said Eren Niazi, CEO and Founder of Open Source Systems. "With 2 discrete servers in each system, you can run 2 separate operating systems within one form factor. This is especially beneficial for our customers who need to run Windows® and Linux® concurrently. You also have the ability for multiple RAID configurations, all while significantly reducing power consumption and ambient noise."

"We are very excited about the performance enhancements of the Next-Generation AMD Opteron processor, with energy-efficient DDR2 memory and outstanding AMD PowerNow!™ technology capabilities," said Jared Giles, Director of Marketing, Open Source Systems. "With planned seamless upgradeability to quad-core made possible with Socket F (1207), we are able to provide a stable computing platform on a consistent power roadmap, delivering the outstanding performance that our customers deserve."

OSS' new product line will also leverage hardware-enabled virtualization solutions using AMD-Virtualization™ (AMD-VTM) technology to improve business functionality and flexibility. Founded on AMD's innovative Direct Connect Architecture, AMD-V enhances all aspects of x86-based virtualization to increase performance with multi-core capabilities and fast memory access, while simplifying virtualization software and reducing CPU overhead.

OSS' new product line will also leverage hardware-enabled virtualization solutions using AMD-Virtualization™ (AMD-VTM) technology to improve business functionality and flexibility. Founded on AMD's innovative Direct Connect Architecture, AMD-V enhances all aspects of x86-based virtualization to increase performance with multi-core capabilities and fast memory access, while simplifying virtualization software and reducing CPU overhead.

www.opensourcesystems.com



SURFnet Selects Roaring Penguin Software's CanIt-Domain-PRO to Manage spam on their Domains

(Ottawa, Ontario) The fight against spam just got easier today in the Netherlands. Roaring Penguin Software Inc., makers of the acclaimed CanIt line of anti-spam software solutions, announced that SURFnet, the Internet Service Provider for all the Universities in the Netherlands has selected CanIt-Domain-PRO to combat spam on their networks. SURFnet provides a high quality, high-speed network specifically designed for higher education and research institutions. SURFnet provides internet connectivity to 750,000 students, faculty and researchers in over 200 organizations including 28 Universities.

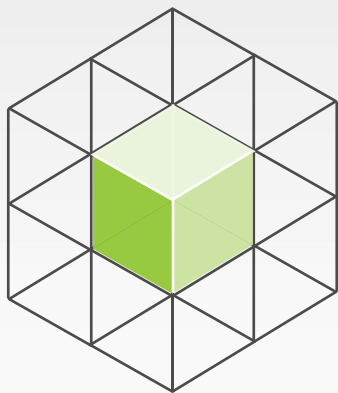
Roaring Penguin Software modified CanIt-Domain-PRO to meet the user functionality and application behaviour that SURFnet expected. These modifications will be fed back into the general release of CanIt-Domain-PRO.

Maurice van den Aker, Project Manager for SURFnet, says, "We recognized early on that Roaring Penguin had an anti-spam solution that was uniquely tailored for Universities. SURFnet had our own solution in-house but we wanted to add end-user functionality among other things. We were looking around for a solution and the University of Twente brought Roaring Penguin to our attention. Purchasing a commercial off-the-shelf solution that had most of the features we wanted just made sense."

David Skoll, President of Roaring Penguin Software, emphasized that Roaring Penguin has extensive experience in the higher education field. "Our North American customers include some of the largest Universities in the United States including the University of Kentucky, Ohio State University, Central Michigan University and Mississippi State University. These customers have unique concerns about how they handle their e-mail, including serious issues around academic freedom and, of course, the ever present problem of tight budgets."

CanIt-Domain-PRO allows a National Research Education Network or a University to delegate control of spam filtering across multiple domains. It is available as a software download, a soft appliance, or as a pre-configured plug-and-play appliance.

www.roaringpenguin.com



Open Management Consortium

Systems Management is Now Open

In May 2006 the Open Management Consortium was announced to help advance the promotion, adoption, development and integration of open source systems /network management software. The founding members of the consortium are Ayamon, Emu Software, Qlusters, Symbiot, Webmin, and Zenoss.

Specific objectives of the Open Management Consortium include:

- Create awareness of open source management tools in the market
- Provide education and resources to help end users make informed decisions regarding open source
- Establish conventions and standards that enable integration and interoperability
- Enable collaboration and coordination on common development projects
- Promote collaborative open source systems management solutions

Open source systems management replaces monolithic vendor lock-in with a modular approach. Pick what you need, customize it to your exact specifications and add to it as your needs change. Because open source products welcome contributions by users, partners and other third parties, they must be standards-based. This standards-based approach facilitates interoperability between open source solutions across the systems management life cycle.



www.open-management.org



~~I CAN'T GET MY DATA OUT OF OUR CRM SYSTEM.
CUSTOMIZING OUR CRM SYSTEM IS OFF-LIMITS.
OUR CRM SYSTEM IS SO INFLEXIBLE.~~

SugarCRM® solves the problems that proprietary CRM solutions created.
Gain control over your business critical data.

Start using the most powerful commercial open source CRM solution for managing sales, marketing, support and group collaboration activities today.

Get hooked on SugarCRM. Visit www.sugarcrm.com/eosm or call +1 408.454.6941.

Commercial Open Source Customer Relationship Management



SUGARCRM®
COMMERCIAL OPEN SOURCE
www.sugarcrm.com • www.sugarforge.org
1.87 SUGARCRM • +1.408.454.6941