

**Linux Cluster**

24/7 Support & Hardware Guarantee-  
Get Managed Solutions & 0%  
Downtime

**Get The Career You Want**

Free Information Pack From The  
Nations Premier Online College. aff

**Virtual Dedicated Servers**

Multiple Virtual Environments per  
account. From \$19.95/month

**Veritas Cluster Server**

Provides high availability for up to 32  
node clusters.



Logged in as maxbaker

[Preferences](#) | [Logout](#)
 
[Home](#) [Topics](#) [Trails](#) [Subscribe](#) [Back Issues](#) [Contact](#) [Contribute](#) [Advertise](#) [Register](#)

## The State of the Penguin: IT

Written by Sam Williams

Tuesday, 13 September 2005

In the summer of 2004, Richard Silver faced a hacking dilemma, literally.

In the course of a converting a large, open air workspace to cubicles, a construction crew at Silver's East Alabama Medical Center had come up with a simple, albeit blunt, solution to a network wiring problem: instead of saving and labeling the pre-existing cables, workers simply chopped the entire trunk of incoming line at the ceiling and pulled new wire, even though keeping the existing wire would have added barely an hour's work to construction. Silver, the medical center's senior networking engineer, was then faced with the daunting task of hooking up 60 new employees with a similar number of dead lines still attached to the network. "I had no way of knowing if a port was off or if it had been cut off," says Silver, looking back.

Fortunately, Silver had NetDisco (<http://netdisco.org/>), an open source port monitoring tool. Taking advantage of the NetDisco's built in archive feature, Silver was able to identify the ports that had shown up on network scans only a few weeks before. Writing those ports off as permanently "dark," he was able to identify and label each new port, the machine it was attached to, and its relative location on the floor.

Not surprisingly, Silver has since become one of NetDisco's most vocal champions on the Internet. "It's really slick," Silver says. "When I first got here, we didn't have any type of port control. I was spending a lot of time with people walking into my office and say, 'Hey can you tell me where this device is plugged in?'"

### Whimsical in Name, Purposeful at Heart

NetDisco, SNORT, Nagios, Rocks. The package names seem whimsical even by open source standards, but for today's IT administrators, those very tools and others are the collective answer to a very serious question: in an era of spiraling network complexity, how do you keep tabs on each and every machine? From intrusion detection (SNORT) to server monitoring (Nagios, OpenNMS), from router traffic (MRTG) to automated configuration tools (cfengine, Rocks), the so-called "stack" of free IT management tools is growing at a speed that evokes memories of Linux' s and FreeBSD' s co-emergence a decade ago.

"It's all about visibility," says Will Winkelstein, Vice President of Marketing for GroundWork, an Emeryville, California company that sells a package of open source monitoring tools dubbed GroundWork Monitor. Bundling Nagios, MRTG, and the open source database MySQL to facilitate trend tracking, the company markets an "IT dashboard" for businesses that leverage Linux to build sophisticated clusters and multi-tiered networks from commodity components.


"Right now, a medium-sized company has a network that's as complicated and robust as any Fortune 500 network would have been five years ago," Winkelstein says.

But the rise in scale and complexity hasn't been matched by a proportional raise in IT budgets. Commercial network management suites such as Hewlett Packard's OpenView or IBM's Tivoli currently require hefty licensing fees, placing the products well outside the budgets of medium-sized companies. Even large companies, eager to swap out high-end



**FREE NEWSLETTER**  
A roundup of news, opinion,  
and targeted industry analysis  
delivered to your inbox weekly

▼ ADVERTISEMENT


 ▼ ADVERTISEMENT  
▼ ADVERTISEMENT

▼ ADVERTISEMENT

▼ ADVERTISEMENT

<a href="#">Back Issues</a>	<a href="#">Linux and the Law</a>
<a href="#">Booting Up</a>	<a href="#">Letters to the Editor</a>
<a href="#">Cluster Rant</a>	<a href="#">MySQL</a>
<a href="#">Do It Yourself</a>	<a href="#">News</a>
<a href="#">Extreme Linux</a>	<a href="#">Perl of Wisdom</a>
<a href="#">Guru Guidance</a>	<a href="#">Power Tools</a>
<a href="#">Hardware</a>	<a href="#">Product Picks</a>
<a href="#">Reviews</a>	<a href="#">Shutdown</a>
<a href="#">Interviews</a>	<a href="#">Software</a>
<a href="#">Kernel</a>	<a href="#">Reviews</a>
<a href="#">Development</a>	<a href="#">System</a>
<a href="#">The LAMP Post</a>	
<a href="#">Java Matters</a>	

Unix systems for low-cost Linux clusters, are turning to open source for help.

The end result of this collective migration has been a surge of development and feedback. "In four years, we've caught up with HP's OpenView, which has been around for fifteen years," boasts Tarus Balog, president of the OpenNMS Group, a North Carolina startup that oversees the development and marketing of the server monitoring tool OpenNMS.

While some engineers might dispute Balog's assertion that open source tools have "caught up" with anything, a quick review of Sourceforge (<http://www.sourceforge.net/>) statistics reveals heavy interest. NetDisco, Nagios, and OpenNMS rank in the 96th, 95th, and 86th percentiles, respectively. Nagios, a project first launched under the name Netsaint in 1999 by Minnesota programmer Ethan Galstad, has already reached its 2.0 milestone, averaging more than 500 downloads a day, according to recent Sourceforge metrics. (For more about Nagios 2.0, see the feature story "The Watcher Knows" in last month's Linux Magazine).

In a recent interview posted by the European open source group FOSDEM (<http://www.fosdem.org/>), Galstad summed up the pace of Nagios development succinctly: "When I first released the code, I thought I'd be absolutely thrilled if twenty or so individuals or organizations used it," Galstad said. "As it turned out, releasing the code triggered something that spiraled out of control."

## Treating IT Like a Business

One of those downloaders is Myke Place, systems administrator for Salt Lake City, Utah Internet service provider (ISP) XMission. Like most Linux network administrators, Place used to write shell scripts to ping individual network nodes and send him a notification whenever a server went down. The scripts worked well for a small number of machines, but when the network grew to 145 nodes, connecting a failure notice with the faulty machine became rather difficult.

"With older[ management] systems, you might lose a router with 20 machines behind it," he says. "And instead of getting that one page about the router, you got 21, one for every device that seemingly failed. You'd get inundated with error messages and still have no idea where the problem occurred."

Now, with the help of Nagios and MRTG, Place can set server and host dependencies and prioritize the error messages that make it to his pager. In addition to separating major problems from the run of the mill "404: Server not found" errors, Nagios has made it easier to spot bottlenecks in network traffic and optimize overall server latency.

Because both Nagios and MRTG have web interfaces, Place has also been able to demonstrate improvements in network performance in a series of tidy charts and graphs. The graphs go out to top-level executives and customers alike. "We wanted to give our customers access to past performance stats and let them decide for themselves whether we're doing a good enough job," says Place.

Place's and his company's approach isn't unique either. GroundWork's Winkelstein says his company's target customer is the systems administrator eager to justify his department's budget to top level corporate managers. "It's the whole notion of treating IT like a business instead of a cost center," says Winkelstein. "If you've got clear data that shows how your network is affecting the bottom line, it makes it that much easier to justify your existence."


Founded just two years ago, Groundwork is one of a handful of companies seeking to capitalize on the growing demand for open source tools and management expertise. Sam Lamonica, IT director for the Foster City, California construction management firm Rudolph& Sletten, credits GroundWork for adding "stability" to the open source IT management tools sector. In the past, Lamonica used Unicenter —a network management tool marketed by Computer Associates —and unsupported versions of Nagios to keep on top of sprawling networks. Now he uses GroundWork Monitor to keep track of his company's 150 servers and workstations.

"As much as I'd love my team to be totally ramped up on open source solutions, they have their own jobs to deal with," Lamonica says. "Being able to bring in an organization like GroundWork, who already has the knowledge and skills to implement a solution and to

Programming

Tech Support

In the Trenches



**CMP Research: Read The Open Enterprise: Mandate for Success and ensure your IT future.**

**Read about the Monarch Empro Enterprise 4-CPU, 8-processor 3U rack server**

train my team in the background, is very helpful."

Other administrators, however, prize the lingering informality of a development community still dominated by hobbyists. Jerry Christopher, a Unix systems administrator at Applied Micro Circuits Corporation in San Diego, uses cfengine to automatically configure his company's network and OpenNMS to keep track of individual machine performance. Having faith in each tool takes some getting used to, he says, but when problems emerge, it isn't too hard to get your questions answered by the very person who wrote the software.

"I once had the opportunity to invite [cfengine creator] Mark Burgess to this company," says Christopher. "He was coming to a LISA conference in San Diego, so I shot him an email and said, 'I'm using your tool everyday and if you ever want to stop by for a beer, you're more than welcome.' Burgess ended up stopping by."

The memory calls to mind the days when it wasn't that unusual to have Linus Torvalds show up at your local Linux user group meeting, and it still impresses Christopher. "I'm not sure that would have happened with HP OpenView," he says.

## The Gory Glory of Success

Then again, for project managers like Max Baker, a Xilinx engineer who created NetDisco during his days as an undergraduate at the University of California at Santa Cruz, the scalability issues that distance a 100-server network from a 10-server network are similar to the scalability issues distancing tool integration from toolmaking. Baker admits to managing NetDisco as "kind of a hobby" and says the project hit a critical mass of users and contributors a little more than a year ago. Thanks to increased feedback, what began as largely a passive attempt to find individual users and devices on a network has become a sophisticated configuration and polling tool for network managers.

But when it comes to expanding NetDisco's features, Baker has to admit that he doesn't have the time. "There's a bunch of tools we could use, tying them together all in one," Baker says. "That's a pretty big undertaking, though. You'd need almost a commercial entity to do that."

The OpenNMS Group's Balog agrees. Balog says he became the "chief steward" of OpenNMS when his former company, Oculan, decided to exit the open source software realm. Rather than let the community built up around the company's Java-based management tool wither and die, Balog decided to launch a company dedicated solely to OpenNMS development and support. In 2004, he merged with another local software company. Balog's current company has 35 enterprise customers and offers a variety of services, everything from remote monitoring to second- and third-level support.

"We scan services on the network like a user," says Balog, describing a sample service. "That means Web server lookup, DNS lookup. We also collect data, monitoring up to 145,000 values every five minutes. The 'third half' we do is notifications and event management: send an alert to the systems administrator, then a page, and then page his boss."

Most of Balog's clients have networks with 200 to 300 devices, a size that still surprises him in retrospect. "When we started this, we thought our average customer would be 'Bubba's Insurance,' some small company that couldn't afford OpenView. But it turns out a lot of our customers are very large corporations that pay a quarter million dollars a year in support fees."

One such customer is Rackspace, the San Antonio, Texas hosting services provider. Eric Evans, a Rackspace senior systems engineer and OpenNMS contributor, says his company uses OpenNMS both as an internal network monitoring tool and as a platform for customer service. For the latter, OpenNMS tracks up to 900 nodes and 60,000 interfaces. Depending on the customer's service level agreement, Rackspace engineers monitor customer servers for uptime and intervene if server or application performance falls below a pre-set threshold.

"If a box goes down, we'll always look into it," says Evans. "But sometimes customers want us to respond to things that don't fall into the category of hardware failure. In such cases, we're passing events to our customer interface, which is integrated with our own internal ticketing system. Our engineers get the ticket, while the customer gets an email letting them know we've responded."

Integrating such service into the existing Rackspace network has been challenging. One reason the company opted for OpenNMS, Evans says, is because the software's open source licensing and background offered a desirable level of flexibility.

"There was pretty much no monitoring solution that was going to work for us out of the box," Evans says. "[With OpenNMS] We felt we had a little more control over what we could do with it."

Such user control is one reason why companies who once based their business strategies on enterprise Linux service are staking a claim in the tools realm instead. A chief example of this migration is Levanta, a San Francisco company that sells configuration management, provisioning, and software deployment tools that allow an IT manager to treat an entire server farm or corporate datacenter as a single, programmable machine. The company is the direct corporate descendant of Linuxcare, the first significant Linux company to attempt a pure service business model. That bold move earned Linuxcare plenty of attention in the brief venture capital heyday of 1998-1999, but it gave the company little to fall back on once main street investors began souring on Linux startups and the software sector in general.

Nevertheless, Linuxcare continued to plug away, developing both the in-house tools and human expertise to help the growing number of customers deal with Linux as an enterprise-grade operating system. By 2003, the company's executives sensed a shift in market demand and repositioned Linuxcare as a tools vendor first. "There was a point in which management decided the best course was the productization of some of the services Linuxcare was offering," Dennis says. To drive home this rebranding effort, the company adopted the name of its main IT product, Levanta, in May of 2004.

Like the offshore outsourcing trend of the same time period, the increased desire for open source management tools sensed by Linuxcare management was a reflection of the harsh economic climate. Companies that might have signed a service contract without batting an eye in 1999 were suddenly looking to slash overhead. In such environments, tools offered a chance to minimize the need for human talent or, at the very least, direct that talent to hard problems beyond the scope of automation.

Dennis, himself a veteran of Hewlett Packard's Openview division, also credits Linux itself for sparking the change. Once limited only to certain corners of the corporate IT infrastructure, the operating system's growing robustness between 2001 and 2003 set in motion a cycle of higher expectations. Executives impressed by the operating system's ability to handle mission critical tasks began demanding the same performance from the tool used to monitor and visualize network performance.

"Remember these people are coming out of Solaris, HP-UX, and AIX environments," Dennis says. "They're more concerned about ecosystems, tools and service. They want something similar to what they're been using for well over a decade now."

AMCC's Christopher echoes that sentiment. "There's a mindset that goes into[ running a network]," he says. "You really need to stop thinking of these individual systems that need to be administered and start thinking of them as one larger system."

## Gentlemen, Start Your cfengine!

Hence the growing popularity of cfengine open source configuration tools. In a recent O'Reilly & Associates survey naming the top five open source automation tools, veteran systems administrator Aeleen Frisch, gave cfengine the top nod over Nagios. "Cfengine is a wonderful tool," wrote Frisch. "cfengine can automatically bring one or a large number of systems into line with each one's individually-defined configuration specifications."

Such work seems simple on paper but is in fact extremely difficult. Christopher says one reason a network administrator might still give cfengine the nod over other, fast-maturing monitoring tools is because, for many administrators, the biggest roadblocks occur during deployment and configuration.

"It's huge," he says of cfengine's impact on deployment. "There's only myself and another Unix admin here for all these systems. cfengine's saved us countless hours, and lets us make changes rapidly and efficiently across hundreds of systems."

Another popular deployment tool is Rocks, a global management tool for cluster-based

supercomputers developed by the San Diego Supercomputer Center in 2000. (For more information about Rocks, see the May 2005 "Extreme Linux" column, available online at <http://www.linux-mag/2005-05/>.) Mason Katz, group leader for cluster development at the center, says the first version was "embarrassingly terrible," but because cluster administrators had few other tools to rely on, it quickly gained users. Such use has prompted a flurry of feedback and improvements, quickly elevating the tool to respectability. He notes the recent integration of Sun Grid Engine, an open source management tool developed by Sun Microsystems. "We were able to support Sun Grid Engine on[ AMD] Opteron[ chips] even before the Sun Grid engine folks were ready to support it," Katz says.

Katz credits the hack to Scalable Systems (<http://www.scalablesystems.com/>), a Singapore company that helps corporations and academic institutions build their own supercomputing clusters using commodity-priced hardware. With the help of Rocks' automated configuration shortcuts, says Scalable Systems president Lawrence Liew, a software tuning process that once took three to five days to complete now takes only five hours. In addition to the Sun Grid Engine, Liew and the Scalable Systems team have added PVFS (<http://www.pvfs.org/>), a parallel, virtual file system tool co-developed by Clemson University and the computer science division of the Argonne National Laboratory.

"Because we use Rocks to generate revenue, its only right that we contribute back to its development," says Liew. Indeed, Scalable Systems has poured so much back into the Rocks codebase that Mason Katz, group leader for cluster development at the San Diego Supercomputer Center credits them as a co-developer of the current software.

"A lot of groups pop up and say they're going to start a cluster distribution only to find it's unbelievably difficult to build a piece of software that works for everybody," Katz says. "The ability for us to integrate other open source pieces of software has been of tremendous value."

## "The Open Source OpenView"

For the moment, such integration has involved minimal overlaps. While it's true that toolkits like Groundwork Monitor and OpenNMS bundle many of the same services—server monitoring, router monitoring, graphical reporting—the field remains largely open for a company to become the OpenView of open source.

That fact, says Winkelstein, is what prompted Groundwork' founders to break away from the world of proprietary monitoring tools and services and stake a claim in the emerging open source frontier. That the company recently gained the backing of Silicon Valley venture capital firm Canaan Partners is yet another sign of the opportunity many technology watchers see.

For the moment, Winkelstein sees big time vendors such Hewlett Packard as the competition, not other open source startups. "The two factors at play are cost and flexibility," he says. "With open source, you essentially can build the system from the ground up. That means you end up with a leaner, more flexible system better suited to what you want to do."

In offering his own assessment of Groundwork, Rudolph and Sletten's Lamonica echoes those same exact points in quick order. "As an open source solution, [Groundwork Monitor] is brilliant for us—its price and features are exceptional," he says. "It doesn't require us to have the cumbersome administrative burden you tend to incur with an OpenView or Tivoli and, frankly, it's a really elegant view. It's a good fit for a smaller IS team, which is exactly what we are."

Sam Williams covers business and software technology for a number of publications. He is the author of two books, "Arguing A.I." and "Free as in Freedom: Richard Stallman's Free Software Crusade."