

INTERNET WEEK

www.internetweek.com

CLUSTERING

► CONTINUED FROM PAGE 43

understand the business benefits of each.

Online advertising leader DoubleClick Inc. is a perfect example of a company that custom-tailors cluster architectures to meet specific business objectives. To serve up its current rate of 15.1 billion ads per month in real time—which requires sustained support for about 12,000 connections per second—the service provider needs significant traffic-handling scalability. The company also must scale its processing capabilities, since it has to provide comprehensive reports on ad activities to its customers throughout the day. And reliability is clearly a must, since revenue depends on ad delivery.

"We're committed to serving the right ad to the right person at the right time," says DoubleClick vice president of MIS and operations Robert Linsky. "So scalability and reliability have to be givens."

As Linsky points out, each of these requirements is fulfilled by a different aspect of his company's network-and-systems architecture. Clustering helps them scale the computing power they need to process reports on ad activity.

"When we pull the log information from the ad servers we have distributed across our network, we end up with this huge file—to the tune of about 40 GB of data," he explains. "So I need a lot of CPU power to crunch it fast and get it to our customers."

That power is generated by 28 Windows NT boxes joined into a cluster using Giganet Inc.'s cLAN technology. Giganet supplies hardware and software that inter-

vidual CPU in this type of configuration has far less impact on the overall performance of the system than it would in a single server," he says. Third, the cluster lets Linsky more easily take individual processors offline for routine software maintenance and upgrades. "We can install a new piece of software on one of the servers without jeopardizing the entire system," Linsky says. "If there's a problem, we can just take it offline again and roll it back."

Finally, there's the issue of piecemeal scalability—which also is not something a big-ticket multiprocessor box can offer. "We can take our capacity up in whatever increments we need to," he says. "In fact, in a little while we'll move up to 32 machines."

DoubleClick's centralized reporting cluster stands in stark contrast with its ad delivery architecture, which relies on 320 servers distributed worldwide across 21 data centers. The servers are linked to the Net via multiple ISP connections and load-balanced using Resonate Inc.'s Global Dispatch and Central Dispatch. Borrowing a term from the storage industry, Linsky refers to this as a "striped" content-serving architecture. "We can lose a data center, an ISP or even a whole geographic area and still be able to serve ads," he says. "But that requires a completely different approach than what we're doing with our back-end reporting application cluster."

Nonstop Service

For Trade.com, clustering is an indispensable component of a larger nonstop service strategy. The financial portal delivers a wide range of services to financial institutions all over the world, including securities transactions, currency exchange, and mergers and

Linux Clustering Extends Trend

LINUX DELIVERS LOTS OF COMPUTING POWER ON COMMODITY Intel processors, and it's especially popular with Net devotees. So could Linux turn out to be the OS of choice for dot.com server clustering?

Plenty of vendors think so. Network Engines recently started to ship its XEngine Linux cluster. And, Linux systems leaders VA Research Inc. and TurboLinux Inc. are also shipping clustering solutions based on Linux.

Some prominent users are buying in: The University of New Mexico built a 128-server cluster using technology from Alta Technology Corp., which runs the Red Hat Linux OS on Pentium II processors. Nicknamed "Roadrunner," the National Science Foundation-sponsored project gives users all over the world access to major computing power via the Internet.

"We're interested in using clusters to solve problems that would overwhelm an individual workstation," says David Bader, an assistant professor of electrical and computer engineering who heads up the project.

"It may be that you need more capacity, or that time is a critical factor in what you're doing," Bader adds. Such applications include highly complex particle physics equations and visualization of huge amounts of data generated by space observation equipment.

"With a Linux cluster, you can write an application on your own workstation and then easily transfer it over to our machine," Bader explains. "With some other large-scale computing platform, you'd either have to have access to it for development, or you'd have to worry about porting it and modifying it after you developed it on your own machine."

Linux clustering is not only for the rarified world of academia. Building clusters with Linux servers offers a powerful bottom line case for business users:

"The two machines I'm using in my TurboLinux cluster cost me \$1,100 apiece," boasts Michael Kuehl, president of Digital Facilities Management Inc., a Web hosting and Internet development firm. "To achieve the same results with a commercial Unix product could cost as much as \$70,000." •

—Lenny Liebmann

David Bader



"We're interested in using clusters to solve problems that would overwhelm an individual workstation."

Photo by Liaison/Cary Hertz

says. "But we're too busy to worry about all transactions. The mirrored systems were



Click here to download
FREE White Paper!



INTERNETWEEK.com
TECHNOLOGY THAT CONNECTS THE ENTERPRISE

Part of the
TechWeb Business Technology Network

[HOME](#) | [BREAKING NEWS](#) | [SUPPLY CHAIN](#) | [WEB DEVELOPMENT](#) | [SECURITY](#) | [IT SERVICES](#) | [ALL STORIES](#)

Enter term(s)

InternetWeek

TechWeb

Search

[Advanced Search](#)

Free Newsletter

Sign up for the FREE InternetWeek NewsBreak e-mail newsletter!
[Subscribe](#)

Focal Points

(sponsored links)

- Business Integration
- Routers & Switches
- VOIP
- Windows Server 2003

Resources

- Home
- Features
- About InternetWeek.com
- Contact Us
- E-Mail Newsletter
- Tech Library
- TechCareers
- Privacy Statement

Resource Centers

Virtual Private Networks (VPNs)

TechWeb Sites

- InformationWeek
- InternetWeek
- Network Computing
- Network Magazine
- Financial Technology Network
- Bank Systems & Technology
- Insurance & Technology
- Wall Street & Technology

SPONSORED BY



FR

Informa
Edit
Persp
Web

Busi
Intelli
Choos
Best V
Gain Ins
Your B

View
On-D
TechW
NO

[View](#)

IN DEPTH

September 20, 1999

Linux Clustering Extends Trend

By **LENNY LIEBMANN**

Linux delivers lots of computing power on commodity Intel processors, and it's especially popular with Net devotees. So could Linux turn out to be the OS of choice for dot.com server clustering?

Plenty of vendors think so. Network Engines recently started to ship its XEngine Linux cluster. And, Linux systems leaders VA Research Inc. and TurboLinux Inc. are also shipping clustering solutions based on Linux.

Some prominent users are buying in: The University of New Mexico built a 128-server cluster using technology from Alta Technology Corp., which runs the Red Hat Linux OS on Pentium II processors. Nicknamed "Roadrunner," the National Science Foundation-sponsored project gives users all over the world access to major computing power via the Internet.

"We're interested in using clusters to solve problems that would overwhelm an individual workstation," says David Bader, an assistant professor of electrical and computer engineering who heads up the project.

"It may be that you need more capacity, or that time is a critical factor in what you're doing," Bader adds. Such applications include highly complex particle physics equations and visualization of huge amounts of data generated by space observation equipment.

"With a Linux cluster, you can write an application on your own workstation and then easily transfer it over to our machine," Bader explains. "With some other large-scale computing platform, you'd either have to have access to it for development, or you'd have to worry about porting it and modifying it after you developed it on your own machine."

Linux clustering is not only for the rarified world of academia. Building clusters with Linux servers offers a powerful bottom line case for business users:

"The two machines I'm using in my TurboLinux cluster cost me \$1,100 apiece," boasts Michael Kuehl, president of Digital Facilities Management Inc., a Web hosting and Internet development firm. "To achieve the same results with a commercial Unix product could cost as much as \$70,000."



Tailoring clusters to meet business apps keeps dot.com Web sites running

[Link to main story](#)

Links to previous In Depth stories:

[Storage Area Networks: A Bet Worth Taking](#)

[Customer Relationship Management: Satisfaction Guaranteed](#)

[Converged Networks: The Bumpy Road To IP](#)

[High-Speed Bandwidth: Terabit Safety Net](#)

[Call Center Breakthrough](#)

[For additional stories](#)