


SUNEXPERT

The Server/Workstation Magazine

INTRANETS STILL NOT READY FOR VIDEO



Premiere: WebServer Magazine Supplement



IF OUR MEMORY IS GOOD ENOUGH
FOR THE PEOPLE WHO MAKE WORKSTATIONS
ISN'T IT GOOD ENOUGH
FOR YOU?



Kingston MEMORY



ARE YOU PARTICULAR when it comes to workstation memory? So are we. The fact is, Kingston® workstation memory is virtually the same as the workstation manufacturers'. But it costs substantially less. Kingston engineers begin with specifications set by the workstation manufacturers. And they optimize the design of every module to suit the system www.kingston.com/ws

in which it will be used. Want proof? The leading workstation manufacturers cover Kingston memory in their on-site service agreements. Now, are you ready to save a fortune on quality workstation memory? Just visit our Web site or call (800) 835-2545 for the name of your nearest Kingston dealer.

Kingston
TECHNOLOGY

COMPUTING WITHOUT LIMITS.™



Kingston Technology Company, 17600 Newhope Street, Fountain Valley, CA 92708 USA, (714) 435-2667, Fax (714) 435-2618. © 1998 Kingston Technology Company. All rights reserved. Computing Without Limits is a trademark of Kingston Technology Company. All other trademarks and registered trademarks are the property of their respective owners.



Having Trouble Accessing Your Headless Servers?

The SystemConsoleSwitch gives you access to and control of multiple headless servers, from one easy to manage location. Consolidate your server room while saving space and energy by eliminating excess terminals. Access to your headless servers is not limited to the server room - use the SystemConsoleSwitch's modem dial-up or telnet to your servers for off-site access.

More than consolidation, the SystemConsoleSwitch provides you with a contingency that keeps you in control of your systems when devices crash. Alternative access to your servers via the console port ensures that you will always have a contingency for control. And while alternative access is a great tool, it would be incomplete without our audit trail, 32k for each port, telling you why your devices crashed.

SystemConsoleSwitch is a SPARC Verified solution.

800 871-9838



**Control
Consolidation
Contingency**




Lightwave Communications, Inc.

261 Pepe's Farm Road • Milford, CT 06460 • fax: (203) 874-0157 • sales@lightwavecom.com
www.lightwavecom.com/sun

SUNEXPERT

The Server/Workstation Magazine

CONTENTS



COVER BY
CHRIS BUTLER

Feature

56 Intranets Still Not Ready for Video

Graphics have come to corporate networks, thanks to large investments in backbone bandwidth, but most company network infrastructures still won't support the next frontier—video applications.

by Paul Korzeniowski

Product Review

64 Cross-Platform Storage for the Masses

This month, we put TAS 5.2—a storage solution that is winning praise from both users and vendors alike—through its paces and, in the final analysis, it's definitely worth a look.

by Ian Westmacott, Technical Editor

News

- 6 Includes:** The Business Case for Java, NC Vendors Court Developers, SP Goes Wildcat, Sun and HP Skirmish Over Workstation Pricing.

Columns

- 16 Ask Mr. Protocol** by Michael O'Brien
Your Basic Standard Mr. Protocol
Mr. P. gives a guided tour of the somewhat bureaucratic Internet standards track.
- 24 UNIX Basics** by Peter Collinson
make: Mastermind of the Update
Using a standard UNIX tool, *make*, to automate the gargantuan task of updating Web pages.
- 32 I/Opener** by Richard Morin
A Special Kind of Freeware
The big names in free software hold a summit to deliberate the future of the freeware community.
- 35 Systems Administration** by S. Lee Henry
Proper Searching
In trying to solve a problem on one of her NT clients, our sys-admin discovers two clever search tools: *agrep* and Glimpse.
- 37 NTegration** by Aileen Frisch
Sharing File Systems, Part 2
In this second column on Windows NT file systems, we consider Samba, a facility for sharing NT resources with UNIX.
- 40 Q&AIX** by Jim Fox
Surmounting Unmount Woes
You can't unmount a file system if any of the files are open. But how do you find those open files?
- 44 Datagrams** by John S. Quarterman
Processing Masses of Mail
procmail is such an essential tool, it deserves a column to itself.
- 48 AIXtensions** by Jim DeRoest
A Picture Worth a Million Bits
Digital vs. film-based cameras: Which you choose depends on what you plan to do with the photos.
- 52 Work** by Jeffreys Copeland and Haemer
We Argue with Our Code
Don't just take them for granted, get the lowdown on command-line arguments.

WebServer

For Managers of World Wide Web Sites

M A G A Z I N E

Visit our Web site at
<http://webserver.cpg.com>

68 UNIX vs. NT Myth

Ian Westmacott, Technical Editor

Ever since its launch, Windows NT has been held up against the venerable tried-and-true UNIX. Both sides have avid proponents and the popular press has been awash in the UNIX vs. NT debate.

73 Communicating without a Net

Simson L. Garfinkel, Contributing Editor

The promise of wireless data communications is simple: free bandwidth. Just buy a pair of wireless Ethernet bridges, put antennas on a couple of roofs, and you'll have the equivalent of a T1 data circuit between two buildings without spending another dime.

78 Clustering



Ian Westmacott, Technical Editor

No matter how big a server is, no matter how fast the processor or how much memory is installed, there are still numerous bottlenecks that affect performance and represent single points of failure. Clustering techniques alleviate these problems for high availability and load balancing of service.

81 A Good Web Server Platform at the Right Price

Ian Darwin

Looking for a Web server platform that is affordable, reliable, secure, scalable and has a list of satisfied users? Then look no further than OpenBSD.



Columns illustrated by
STEPHEN SCHILDBACH

BONUS TO ADVERTISERS OF THE JULY ISSUE:

Siggraph '98, Orlando, FL

BONUS TO ADVERTISERS OF THE AUGUST ISSUE:

Internet Expo, Boston, MA

BONUS TO ADVERTISERS OF THE SEPTEMBER ISSUE:

Comdex Enterprise West, San Francisco, CA
IT Forum, New York, NY
Internet Commerce Expo, Los Angeles, CA

BONUS TO ADVERTISERS OF THE OCTOBER ISSUE:

Fall Internet World, New York, NY



Departments


- 4 Editorial
- 62 Reader Feedback
- 86 New Products
- 89 Server/Workstation Marketplace
- 104 Advertisers' Index

Page 86

SUNEXPERT Magazine (ISSN 1053-9239) is published monthly by Computer Publishing Group, 320 Washington St., Brookline, MA 02146. Telephone (617) 739-7001. Periodicals Postage Rates paid at Boston, MA, and at additional mailing offices. Posted under Canadian IPM #0235873. This publication is free to qualified subscribers as determined by the publisher. Subscription rates are \$60 per year in the United States, and \$95 (surface mail) and \$150 (air mail) outside the United States. Subscription requests can be sent to Circulation Department, SUNEXPERT Magazine, 320 Washington St., Brookline, MA 02146 or electronically mailed to circ@cpg.com.

POSTMASTER, please send all address changes to SUNEXPERT Magazine, Circulation Department, 320 Washington St., Brookline, MA 02146. Please allow 6-8 weeks for change of address. Include your old address as well as new-enclosing, if possible, an address label from a recent issue. All rights reserved. © Copyright 1998, Computer Publishing Group. No part of this publication may be transmitted or reproduced in any form by any means without permission in writing from the publisher.

Material for publication should be sent to the attention of Doug Pryor at the above address or electronically mailed to dpryor@cpg.com. Letters sent to the publication become the property of the publication and are assumed to be intended for publication and may be used so. SUNEXPERT Magazine is not sponsored or endorsed in any way by Sun Microsystems Inc. All information herein is believed to be accurate to the best of our ability.



Video, Shared Storage and Something New

Although words like collisions, dropped packets and Ethernet storms strike terror in the hearts of network managers, there is another word popping up more and more in relation to corporate backbones that has some capacity planners cringing. That word is video.

Makes you cringe just thinking about all those saturated TCP/IP networks. So we decided to ask Paul Korzeniowski to investigate. His report on video and networking appears as this month's cover story, "Intranets Still Not Ready for Video," Page 56. His basic conclusion is that most company network infrastructures still won't support video applications. And no wonder: a 30-second video clip can consist of 25 MB of data. "While there has been plenty of hype about the benefits video offers, corporations have shown little interest in the technology, and it is unclear if or when that will change," says Paul.

This month's issue also contains a review that might impact your network—in a positive way for a change. If your network is anything like ours—a mix of UNIX, Mac OS, Windows 95 and NT—you're always searching for ways to economize and make storage more efficient. Technical editor Ian Westmacott test drives a promising suite of software called TotalNET Advanced Server, or TAS.

TAS allows transparent client access. You install the server—AIX, Solaris, HP-UX, IRIX, or most popular versions of UNIX—and you're done. There's no client-side configuration or software. Now, Macintosh clients (through the Chooser) or 95 and NT clients (through Network Neighborhood) can browse and use file and print services. There are a few snags, but everything works as advertised—you can't say that too often in this business.

There's something new in this issue, too. In the interest of fully covering the UNIX and large system NT business, *SunExpert* will be featuring a new section called WebServer Magazine Supplement. The four articles in this month's installment cover a wide range of issues, from UNIX versus NT to wireless networks, from clustered Web servers to a free Web server OS.

We would very much appreciate your feedback on how we can make this section more useful to you. If you have any suggestions drop them in our suggestion box—editors@cpg.com.

Doug Pryor

Publisher	S. HENRY SACKS shs@cpg.com
Editor-in-Chief	DOUGLAS PRYOR dpryor@cpg.com
Managing Editor	LISA BUCHER lisab@cpg.com
Technical Editors	IAN WESTMACOTT ianw@cpg.com RICHARD MORIN rdm@cpg.com
Contributing Editors	MICHAEL JAY TUCKER SIMSON L. GARFINKEL MARK SEIDEN
Research Editor	MAUREEN MCKEON mm@cpg.com
Staff Editors	ALEXANDRA BARRETT alex@cpg.com PATRICK T. COLEMAN pat@cpg.com SUZANNE HILDRETH sjh@cpg.com
Production Editor	MICHAEL LETTS mletts@cpg.com
Art/Production Director	JOHN W. KELLEY JR. jwk@cpg.com
Senior Designer	JERRY COGLIANO jerry@cpg.com
Designer	BRAD DILLMAN bdillman@cpg.com
Production Assistant	CAMILLE L. DASTOLI cld@cpg.com
Circulation Director	DEBORAH MOORE dm@cpg.com
Promotion Manager	MELANIE DECAROLIS melanie@cpg.com
Accounting Manager	ROBIN YEARWOOD
Office Manager	TINA JACKSON jamal@cpg.com
Business Manager	SUSAN R. SACKS srs@cpg.com

World Wide Web

<http://www.cpg.com>

EDITORIAL OFFICES

320 Washington Street • Brookline, MA 02146
(617) 739-7002



Printed in USA

Western Scientific's
CycloneRAID™:

POWER in the **TOWER.**

CycloneRAID Ultra-2 and CycloneRAID-FC are the trusted RAID solutions for the world's most critical data storage needs, providing unsurpassed speed, performance and reliability.

Western Scientific's CycloneRAID is the solution of choice, from research, seismic data recording, space exploration and real-time data acquisition, to internet services and visual effects production. When you require guaranteed data availability, system uptime and maximum performance, call Western Scientific first. It's the only call you'll need to make.

Up to 1.5TB per controller

Supports RAID Levels 3 and 5

60MB/sec sustained transfer rate with FC

100MB/sec burst rate with FC,
80MB/sec with Ultra-2 Wide SCSI

Configurable up to 70 drives

32MB of segmentable on-board cache memory, expandable to 256MB.

**UNSURPASSED
RELIABILITY
AND PERFORMANCE**

CycloneRAID™
Trusted data storage solutions
from Western Scientific.



Western Scientific, Inc.
9445 Farnham Street, San Diego, CA 92123 USA
Tel: 619-565-6699 Fax: 619-565-6938
Email: info@wsm.com

FOR OUR 1998 CATALOG, CALL
1-800-443-6699
Visit us at <http://www.wsm.com>

**WESTERN
SCIENTIFIC**

©1998 Western Scientific, Inc. All CycloneRAID products and logos are trademarks of Western Scientific, Inc.

The Business Case for Java

Java isn't just for applets anymore. Increasingly, developers use Java to build stand-alone, enterprise-class applications to enable disparate clients, such as Sun Microsystems Inc. workstations, desktop PCs running Microsoft Corp. Windows NT or Java-based Network Computers, to access legacy databases and share applications across a network.

Enterprise Java, or server-side Java, as it's often called, has become a new darling of the industry. In a 1997 survey of U.S.-based mid-size companies, Zona Research, a market research firm based in Redwood City, CA, found that 43% of commercial software development today involves some type of Java project, and that 97% of those surveyed plan to deploy at least one Java application on the server.

The first few months of this year have, in fact, produced a flurry of enterprise Java product announcements. Sun, the leading proponent of Java for the enterprise, has announced the commercial availability of Java JumpStart 1.1, a package of applications for developing and deploying enterprise Java applications. Sun has also released JavaServer Engine, a collection of reusable Java classes for constructing server-based applications.

But the most significant of Sun's announcements so far was the release of the final specifications for Enterprise JavaBeans, unveiled at the JavaONE conference in San Francisco in March. At the same time, 18 other companies, including IBM Corp. and Netscape Communications Corp., announced plans to support the Enterprise JavaBeans specification in their own products.

Enterprise JavaBeans are Java compo-

nents developers can use to construct server applications without having to deal with many low-level system functions, such as session-management and linking to back-end systems. If the Enterprise JavaBeans specification fulfills the expectations of Sun, IBM and the other proponents of server-side Java, it promises to speed the development of large-scale Java applications substantially.

But corporate developers must wait until commercial software vendors begin to ship Enterprise JavaBeans products. At the moment, there is just one



Enterprise JavaBeans product on the market—a spreadsheet engine called Integer (at \$3,000 per server), produced by Athena Design Inc., San Francisco. However, a Sun spokeswoman predicts that once independent software vendors have a final specification to work with, Enterprise JavaBeans products will arrive on the market in a matter of months—by September at the latest.

If the market for Enterprise JavaBeans has yet to take root, the one for Java application servers is entering full bloom.

Application servers hold the application logic of one or more Java programs and communicate with databases and

other legacy systems on the back end and a variety of clients on the front end. An application server typically consists of a Java Virtual Machine (JVM), Open Database Connectivity (ODBC) and/or Java Database Connectivity (JDBC) drivers for connecting to databases and various tools for developing and deploying Java applications. Many also provide network management features, such as load balancing.

In some cases, an application server may also double as a Web server, if it supports HTTP, and may also be capable of running programs in other languages. Netscape's Application Server 2.0, for instance, supports both Java and C++.

There are already about a dozen such application servers on the market, from companies such as Oracle Corp., Redwood Shores, CA; WebLogic Inc., San Francisco; GemStone Systems Inc., Beaverton, OR; SilverStream Software Inc., Burlington, MA; and Novera Inc., also located in Burlington. In

addition, IBM has just released the beta version of its Servlet-Express, a software plug-in that turns a Web server into a Java application server, see <http://www.ibm.com/java/servexp>. And just about every month, yet another Java application server product hits the market. As Ron Rappaport, analyst for Zona Research, says, "It's very chic to have an application server these days. Every company's got one."

But before Java can become a serious contender as an enterprise-ready language, it must show that it can offer the same kind of performance as C/C++. Performance has been improving thanks to better Just-In-Time (JIT) compilers, such as Sun's HotSpot Virtual Machine adaptive compiler, scheduled to be released by the end of the year.

Prism Library Architecture™

Created for growth and capable of adapting to new technology. Prism's industry standard PCI slots allow for a future of high performance PCI adapters. Fibre-channel host adapters. Gigabit network adapters. Single-board computers with embedded backup software. Tape array adapters with mirroring, automatic drive fail-over and more!



Announcing the birth of a new architecture in today's most powerful DLT library

Fibre Ready.
Network Ready.
Server Ready.
Array Ready.

ATL
PRODUCTS

The P1000 Series. A bundle of performance, capacity and fault-tolerance that's incredibly easy-to-use. Up to 30 cartridges and 4 DLT™7000 drives. Rackmount or desk-side enclosures. Hot-swap DLT drives. Security, admin and service from a browser-like touch screen. WebAdmin™ software for library administration anywhere on the world wide web. Prism Inside.

One day all libraries will have Prism Library Architecture.



Quantum
CAPACITY FOR THE EXTRAORDINARY

www.atlp.com 714.774.6900

Circle No. 4

© ATL Products, Inc. DLT is a trademark of Quantum. WebAdmin is a trademark of ATL.

HotSpot VM converts byte code to native code, like other JIT compilers, but does so only on selected, performance-critical parts of the application. Because it converts only part of the application's methods to native code, it has time to do a better job of optimization than a regular JIT-compiling Virtual Machine, according to Sun. Rather than wait for HotSpot VM, however, some developers are taking a different tack. They're opting to compile portions of the code into native code optimized for whatever CPU it happens to be running on.

One such compiler comes from Tower Technology Corp., Austin, TX. TowerJ 2.0 is designed to provide faster performance levels by compiling portions of the Java byte code to native code before runtime—a process called static compilation. At the same time, it leaves other portions of the code to be compiled dynamically. IBM has a similar product, called the High Performance Compiler for Java (the beta version can be downloaded at <http://www.alphaworks.ibm.com/formula.nsf/toolpreview>).

The upshot of using such an approach is that the Java application remains cross-platform for clients accessing it, but must be recompiled before it can be run on a different server platform—a sacrifice that many developers may be willing to make in order to get better performance. Chris Carrigan, analyst for Summit Strategies in Boston, says, "I think you're definitely going to see cross-platform portability become less important, at least for server-side Java development."

In fact, with the number of different versions of Java coming onto the market—including Sun's 100% Pure Java, Microsoft's Windows-optimized Java, Hewlett-Packard Co.'s Java for embedded systems and IBM and Sun's jointly announced JavaOS for Business—the market is already showing signs of choosing functionality over portability.

Or as Greg Vogel, analyst with NationsBanc Montgomery Securities LLC in San Francisco, says: "At the end of the day, customers are simply

looking for solutions to their problems. And if a Java application isn't providing fast performance, that's a problem."—*sjb*

NC Vendors Court Developers

Network computing and Java enthusiasts received a boost when Sun Microsystems Inc. and IBM Corp. declared their intentions to develop and market JavaOS for Business. The strategic announcement reaffirmed both companies' dedication to Network Computers (NCs) by promising to jointly develop a Java operating system specifically for the fledgling market. "We're creating a platform for open standards that will help accelerate adoption of Network Computers and network computing," says Mike Lawrie, general manager of the Network Computing Software Division at IBM.

The agreement between the two companies is designed to hasten the development of Java applications that will be well suited for NCs. Analysts say this show of unity between the companies is a much needed confidence boost for independent software vendors and potential customers. "This wasn't about anything you could get this week. It was more about reassuring the market that as they need more robust versions of the operating systems, they would in fact be there," says Amy Wohl, president of Wohl Associates, a consulting firm based in Narberth, PA. "If you look across the market there is basically a confidence problem."

One example cited for the shaky confidence of this nascent market was the year-and-a-half delay of Sun's own NC, the JavaStation. Sun pushed back the introduction of the product twice during that period. Furthermore, viable alternatives have surfaced in the form of Wyse Technology Inc.'s Winterm products, Citrix Systems Inc.'s Winframe software and the upcoming

release of Microsoft Corp.'s Windows Terminal Server. "When customers are buying into a new market they like to look around and have the reassurance of seeing lots of other people are making the same decision," Wohl says.

To help NCs realize commercial success it has become crucial to persuade developers to create Java applications as opposed to Windows applications. Sun and IBM will offer channel initiatives, technology seminars and educational programs for ISVs and OEMs to encourage this much needed Java development. An OEM version of JavaOS will be released this summer. "We have to give end users and customers specific guidance on how to write for it," says Janpieter Scheerder, president of SunSoft Inc.

The commercial release of JavaOS for Business is scheduled for first-quarter 1999. Sun will also migrate customers using the JavaStation product line from JavaOS for NCs to the

JavaOS for Business software over the next 12 months.

While the NC market is still developing, IBM and Sun are making strong claims as the dominant players—at least in terms of public perception. A March study released by International Data Corp., a market research firm based in Framingham, MA, says IBM, Sun and Hewlett-Packard Co. are "perceived most strongly" with NCs. The study also

indicates that 73% of those who responded purchased NCs to replace PCs in 1997 and 80% bought NCs to reduce the total cost of ownership. Only 25% said Java was a major factor in their decision to purchase a Network Computer, while 34% said it was a non-factor.

"Future purchase plans of NCs show them diffusing slowly through enterprises," says Eileen O'Brien, director of IDC's Enterprise Network Computers research program. "This could change if NC vendors take proactive steps to help users realize what an NC is and why they need one."


Besides encouraging ISVs to develop Java applications, the agreement between

To help NCs realize commercial success it has become crucial to persuade developers to create Java applications.




Imagine a 30 Terabyte Hard Disk Archive Without Spending Terabucks

Imagine the luxury of archiving *everything* to a virtual disk. With AMASS® software you can directly access the enormous storage capacity of an automated library, with the familiarity and ease-of-use of a hard disk. AMASS software makes data stored in optical or tape libraries accessible with the look-and-feel of hard disk — transparent to applications, such as Databases,




Imaging, and Multimedia. AMASS supports the most popular UNIX and Windows NT* platforms. And, you have the freedom to choose the automated library that best meets your specific storage requirements — from 16 different manufacturers such as EMASS, HP, DISC, StorageTek — in capacities from gigabytes to



thousands of terabytes. AMASS is a proven, leading archive software solution, with a successful world wide installed base.

So, expand your archiving abilities, call now for a free AMASS overview book or a 30 day free software trial.



1-800-653-6277 • storage@emass.com • <http://www.emass.com/se>

* Initial AMASS NT release limits supported storage devices. EMASS and AMASS are registered trademarks of EMASS, Inc. © Copyright 1998

Circle No. 34

IBM and Sun also shows unity on the Java front. This was important to Sun following HP's announcement in March that it plans to market its own variant of Java for embedded systems. While Sun and HP deal with products at opposite ends of the spectrum in terms of complexity and size, the HP announcement indicates there is a problem in creating industry unity around Java.

"The HP announcement was damaging," says Marty Gruhn, vice president of Internet business solutions at Summit Strategies Inc. in Boston. "The JavaOS announcement helped mitigate some of that but it didn't solve the whole problem that the industry has with regards to whether or not Java is ever going to be a platform that everyone can stand up and salute."—*ptc*

SP Goes Wildcat


IBM Corp. has unveiled two new nodes, both code-named Wildcat, for the RS/6000 SP, and has announced other hardware and software enhancements for the supercomputer. The April announcements mark the first time the PCI standard has been used in an SP node, and are based on the company's 332-MHz PowerPC 604e microprocessor. The nodes utilize the same microprocessor technology previously delivered in the RS/6000 F50 Server. "We have effectively repackaged [the F50] as a node," says David Gilardi, RS/6000 program manager for IBM business intelligence solutions. "If you look at our history, we pretty much have taken existing servers and made them into nodes."

In general, analysts see this as standard operating procedure for IBM with regard to the SP. "Basically, the SP is kind of the Volkswagon of the computer industry," says Christopher Willard, research director of high performance technologies at International Data Corp., a market research firm in Framingham, MA. "It's the product that keeps on changing but remains the same. It seems every six months IBM will come up with a product improvement for the SP."

The SP houses a collection of processing nodes that come in three types:

Thin, Wide and High. The new nodes are available as either two- or four-way SMP Thin and Wide nodes. The new Wildcat nodes are similar to High nodes in that they are both symmetric multiprocessors. A new four-way Thin node takes up a quarter of the space that would be occupied by an eight-way High node offering similar performance.

For example, IBM's Deep Blue, the supercomputer that defeated chess grand master Gary Kasparov, was a 32-node system. According to IBM, if a one-for-



In a series of announcements, IBM unveiled two new Wildcat nodes and a host of other enhancements for its RS/6000 SP.

one replacement of its nodes were made with new four-way 332-MHz nodes, it would increase the performance capability of Deep Blue from 200 million to more than one billion chess moves per second. "From a performance standpoint that's a significant increase in power," Gilardi says.

In addition to boosting the performance of the SP, IBM has added a new option to its upgrade path for the supercomputer. Typically, the company introduces a server that it later repackages as a node, as it has done with the Wildcat node evolving from the RS/6000 F50. Big Blue altered that plan slightly with the announcement that it plans to introduce an upgrade to the RS/6000 S70 enterprise Web server—a 12-way SMP machine that will allow SP users to

extend out to that server. Basically, IBM will upgrade the S70 so that it can connect to the SP. This connection will allow the SP management tools to reach the S70 as if it was a node.

The ability to connect the S70 and have it function as an SP node will be available by the end of the year, IBM says. Robert Pimienta, systems administrator with Panasonic, a subsidiary of Matsushita Electric Corp. of America, Secaucus, NJ, is already interested in this capability. Panasonic currently runs an order-entry application called System ESS from Industri-Matematik International Corp., based in Stockholm, Sweden, on a UNIX machine from Hewlett-Packard Co. Pimienta describes the program as a "hog" that doesn't fit on any SP node, but would fit nicely on the 12-way S70 with 4 GB of RAM. By migrating the application to the S70, he says he could take advantage of the six-frame SP his company already has. "We'd like to connect it to the SP for the simple reason that we can administer it from one control workstation," Pimienta says.

While a final decision has not yet been made, Pimienta feels the odds are good that Panasonic will go that way. "We haven't decided yet," he says. "There's an 80% chance we'll go in that direction."


This new approach is a look at how IBM will introduce future products. "We're not trying to fit the S70 underneath the frame of the SP," says Gilardi. "We are going to take other new servers and give them the capability to be part of the SP without having to be physically in it." In addition to the new nodes and plans for S70 connectivity, IBM introduced the RS/6000 SP Switch Router Model 16S in April. The new router is designed to allow the movement of large amounts of data between the SP and external networks. The company also introduced AIX Version 4.3.1. This latest release of IBM's UNIX operating system features enhancements in scalability for 32- and 64-bit applications. In addition, IBM says it is also working to improve the performance of Java, Web-based systems management, network security and TCP/IP.—*ptc*

A BGA Heat Sink Should Be Judged By How It Attaches To Your Chip.



Our patented Cool Snap® design is the world's best thermal cooling system for surface-mount and BGA chips. Its revolutionary threaded design provides precision heat-sink contact with your chip, guaranteeing 100% performance. Our Cool Link™ permanent dry-screened interface eliminates thermal grease and other inferior bonding methods no matter how many times you remove or reattach it to your chip. And, because our manufacturing process is as unique as our design, we can provide you with the perfect thermal cooler at the most competitive price. Sun Microsystems, Apple, Motorola, IBM, Digital, Hewlett Packard, Lucent Technologies and Intel are some of the **leading** companies that have discovered the Chip Coolers™ advantage. Call us at 1-800-227-0254, or contact us on the World Wide Web at www.chipcoolers.com about your specific application.

15-50 mm BGA, CBGA and PBGA devices, as well as
21 & 25 mm PowerPC™ processors



Sun Microsystems, Apple, Motorola, IBM, Digital, Hewlett Packard, Lucent Technologies, and Intel are registered trademarks of these respective companies.

Circle No. 5

Sun and HP Skirmish Over Workstation Pricing

In the ongoing struggle to gain—or at least maintain—market share, Sun Microsystems Inc. and Hewlett-Packard Co. have cut prices on their workstations. Following the January introduction of Sun's Ultra 5, Ultra 10 and Ultra 60 PCI-based workstations, HP slashed prices on its Visualize line of technical workstations in late March by 22% to 37%. HP's Visualize Model J2240, which HP's product literature pits against Sun's high-end Ultra 60 workstation, now costs \$35,500 for a base configuration of 1 GB of RAM, a 4 GB hard drive and a 19-inch monitor. The mid-range B180, aimed at users of Sun's Ultra 10, now costs \$11,300 for a base configuration of 128 MB of RAM, 4-GB hard drive and a 19-inch monitor.

As one might expect, the aim of these price reductions is to take market share away from Sun, as well as other workstation vendors, according to HP's product marketing manager for technical workstations, Barry Crume. "We're repricing so we can improve our competitive position," says Crume. "I'd expect other companies to adjust."

And Sun has done just that. In May, the company introduced a higher performance version of its high-end Ultra 60 workstation and instituted price cuts for existing models of the Ultra 60 and mid-range Ultra 10. The Ultra 10, which formerly sold for \$10,395 for a 300-MHz UltraSPARC-III chip, with a 4.3-GB hard drive, 128 MB of RAM and Elite3D graphics card, now costs \$9,195. In third-quarter 1998, Sun plans to unveil a higher-performance version of the Ultra 10 with a 333-MHz chip. Sun dropped the price of its 300-MHz Ultra 60 with a 4-GB hard drive, 128 MB of RAM and Creator3D graphics from \$13,295 to \$11,595. A new, higher-performance model, sporting a 360-MHz UltraSPARC-II chip and a 9.2-GB hard

drive is also available for \$13,595. In addition, Sun reduced its entry-level Ultra 5 workstation from \$2,995 to \$2,495 for a basic configuration with a 270-MHz UltraSPARC-III chip, 64 MB of RAM and a 4.3-GB hard drive.

"We see HP's price reductions as a reaction [to Sun's January introductions of its Ultra 5, 10 and 60 workstations]. This new pricing will just leave HP with more catching up to do," says Robert Novak, group manager for power desktop workstations at Sun.

Greg Weiss, analyst with D.H. Brown Associates Inc., Port Chester, NY, says falling memory prices throughout the industry have helped make it possible for Sun and HP to lower their prices. "RAM prices have dropped by a factor of 16 in the past two years. Disk drive and monitor prices have also fallen," Weiss says.



Sun cuts the price of its 300-MHz Ultra 10 workstation from \$10,395 to \$9,195.

A survey conducted by D.H. Brown in April reports that Digital Equipment Corp., IBM Corp. and Silicon Graphics Inc. have joined Sun and HP in reducing memory prices on workstation RAM and disk

drives. DEC, for instance, has reduced its memory prices by up to 32% and its disk prices by up to 17%. IBM also lowered its disk drive prices by up to 22%.


Despite competition from other vendors, according to D.H. Brown, Sun remains a price/performance leader, tying with DEC for first place in both the SPECint95 and SPECfp95 benchmarks, which measure integer performance in mathematical equations and floating-point performance, respectively. A recent survey of the UNIX workstation market by Dataquest, a research firm based in San Jose, CA, found Sun still firmly in the lead, with 44% of unit shipments and 41.2% of market revenue. HP took second place with 16.8% of unit shipments and 23.1% of market revenue. However, HP still leads in overall workstation shipments—including both Windows NT and UNIX systems—according to International Data Corp., Framingham, MA.—*sjb*

Sun's Latest i-Series Microprocessor


Sun Microsystems has released upgrades to its 64-bit UltraSPARC-III microprocessor line, which was unveiled in January. The 333-MHz UltraSPARC-IIIi, which is targeted at the embedded and entry-level workstation and server markets, is binary compatible with the UltraSPARC family and is offered in both chip and module configurations. According to Sun, the i-Series microprocessors focus on integration and performance.

The first product to take advantage of the 333-MHz microprocessor is Sun's Ultra AXi PCI-based motherboard, designed for OEMs creating high-performance SPARC embedded applications. "It [the Ultra AXi motherboard] is obviously developed with the intention of targeting the PCI-based market," says Andrew Allison, an independent analyst in Carmel, CA, who specializes in the RISC and UNIX markets. "It's a good product that will enable Sun to extend into the lower end of the market and compete with NTs."

"The big news here is that before, you typically had to pay a premium at the systems level to utilize Sun technology," says Jeff Veis, group marketing manager of Sun's OEM platform products. "But coming out with the i-Series and making it available on the board level has changed that."



The 333-MHz UltraSPARC-IIIi comes in both chip and module configurations.



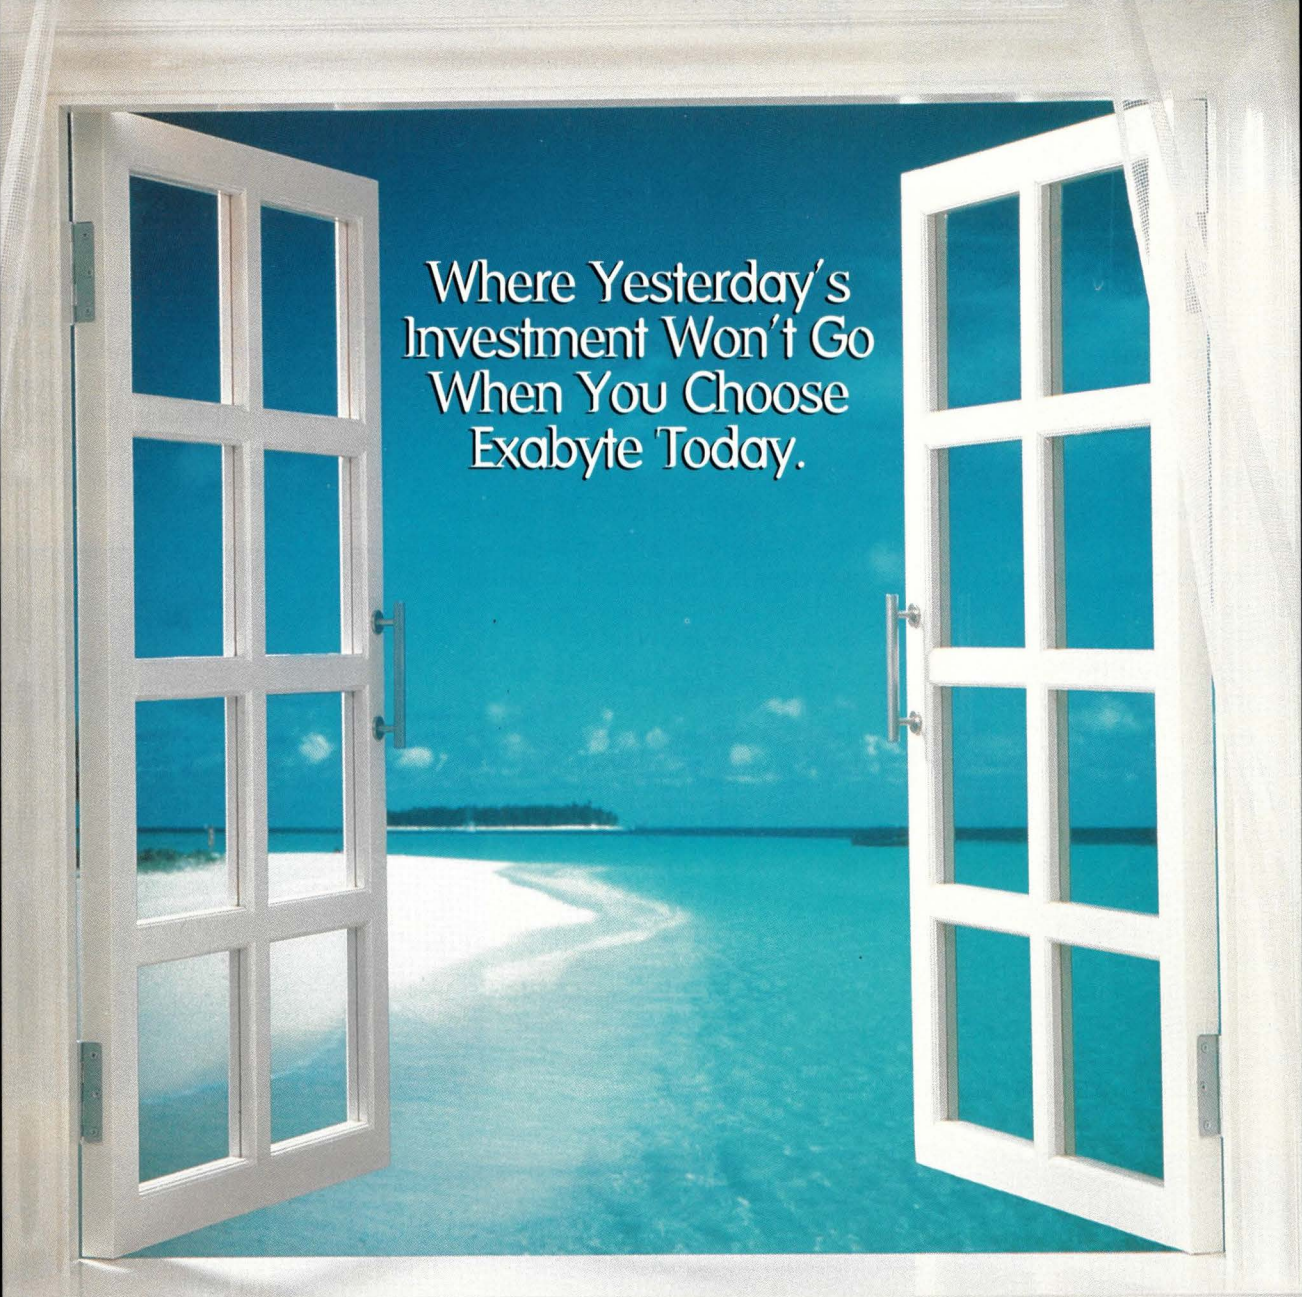
We know your window to the future depends a great deal on where you've been. A tape backup solution can't ignore your investments in platform, software and media; or it's no solution. If you've invested in 8mm, we've got you covered with industry-leading automation

and an aggressive, compatible roadmap in Mammoth technology; if you are using DLTtape™, we're the choice for superior automation there, too. Plus we offer service and support that's second to none. Take a look at Exabyte...you'll love the view.




Exabyte®
every day. every time.

www.exabyte.com



Where Yesterday's
Investment Won't Go
When You Choose
Exabyte Today.

© 1998 Exabyte Corporation. Exabyte is a registered trademark of Exabyte Corporation. All other trade names referenced are property of their respective owners.
Exabyte Corporation, 1685 38th Street, Boulder, Colorado 80301 USA.



The Ultra AXi PCI-based motherboard is the first product to take advantage of Sun's new 333-MHz microprocessor.

Toshiba America Information Systems Inc., Irvine, CA, is currently using the new Ultra AXi motherboard to develop its next-generation cable model access systems.

The Ultra AXi motherboard is available in 300- and 270-MHz versions for \$1,700 and \$1,500, respectively, in quantities of 500 or more. A 333-MHz version is scheduled to be available this month.

When compared with Intel Corp.'s Pentium II 333-MHz chip with a 66-MHz external bus, Sun says the 333-MHz UltraSPARC-III with a 167-MHz internal bus offers 2 MB of external cache to the Pentium II's 512 KB. In addition, the UltraSPARC-III offers a 14.2-SPECint over the Pentium's 12.8-SPECint and a SPECfp of 16.9 compared with a SPECfp of 9.14 for the Pentium II.

The UltraSPARC-III has an integrated 66-MHz PCI and memory controller to maximize throughput and an optional Advanced PCI Bridge to expand the number of PCI devices supported, which is available for \$50 from Sun.

In addition, the UltraSPARC-III takes full advantage of Sun's VIS instruction set for applications in accelerating multimedia, image processing and networking. VIS is a high-performance, robust set of multimedia extensions. Sold in quantities of 1,000 units, the 333-MHz UltraSPARC-III microprocessor costs \$720.—*mm*

Sun Workstation Watch

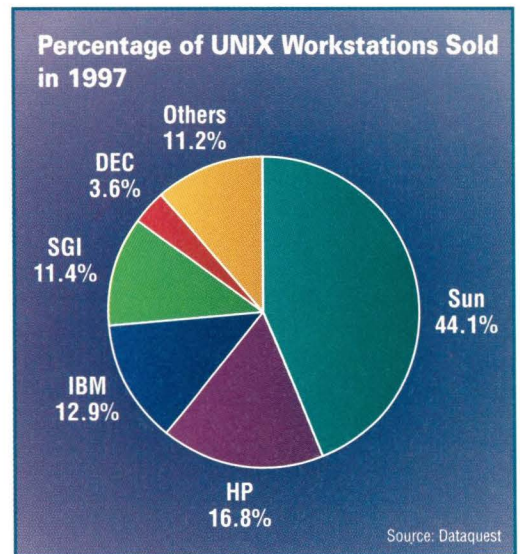
A new, high-end workstation will join Sun Microsystems Inc.'s family of UNIX-based SPARC and UltraSPARC technical workstations. The Ultra 450, which is capable of holding up to four 300-MHz CPUs, 4 GB of memory and 180 GB of hard disk storage, tops off Sun's workstation line and surpasses the previous top-of-the-line Ultra 60 workstation.

The Ultra 450, which will be available in third-quarter 1998, has not yet been priced. But Sun expects a base configuration of the Ultra 450 to help the company win more marketshare on the high-end, while it wages war against Windows NT with its Ultra 5 and Ultra 10 workstations on the low-end.

"As Intel [Corp.] has caught up on the low-end, UNIX RISC vendors have focused on providing better floating-point performance. Sun is shifting its focus to higher-end problems," says Peter ffoulkes director of the Advanced Desktop and Workstation program at Dataquest, a San Jose, CA-based market research firm. "The 450 plays well into that market. It offers the ability to handle more complex problems and larger data sets at the high end."

Targeted at data-intensive scientific and industrial applications, the Ultra 450 can hold up to three Elite3D M6 graphics boards—Sun's high-performance graphics system for 3D visualization and rendering that features six floating-point processors. In addition, the 450 comes with 10 PCI slots, divided among six PCI buses. It boasts 20 internal disk drives, collectively providing a maximum of 180 GB of storage, with four of the 20 hot-pluggable so that users can replace full drives, as needed, without having to power down the system. According to Robert Novak, group manager for Power Desktop Workstations at Sun, that makes the Ultra 450 ideal for applications involving very large files—such as medical imaging, where a few patient scans could fill a single drive.

According to the latest Dataquest figures, Sun still has a healthy share of the UNIX workstation market—44% of unit sales and just over 41% of market revenue. Hewlett-Packard Co. pulls up a distant second with close to 17% of unit sales and just over 23% of market revenue.—*sjb*



WebServer Magazine Back In Print!

Starting with the June 1998 issue of *SunExpert*, editorial from our online magazine will form the basis for a WebServer Supplement. Editorial will include expert discussions on Web security, Web tools, new Internet products and other Web-related topics.

New Products



Sun compatible workstations from Tatung

Tatung now offers Ultra-10 and Ultra-60 Sun compatible workstations featuring PCI I/O busses. Also available are the Ultra-1 and Ultra-2 Sun Compatible systems. The workstations all run on Sun Solaris 2.x and utilize Sun Processors. These units are the cost-effective solution that every Systems Administrator should seriously consider.

800-692-4343
Circle No. 40



Qualstar Tape Libraries

Qualstar designs and manufactures three families of libraries for the popular AIT, DLT, Exabyte 8mm, and DDS tape formats. Various models house 1 to 6 tape drives and from 10 to 144 cartridges. Barcode scanning and I/O Ports are supported in most models. Nineteen models provide combinations of capacity and throughput to match any storage management application requirement. Our software partners provide a full range of backup, concurrent backup, HSM and archival packages for virtually every platform.


800-692-4343
Circle No. 41



Dataram Memory Upgrades

Dataram, a leading provider of Gigabyte memory upgrades for UNIX workstations, now offers memory upgrades for the Sun Ultra-5 and Ultra-10 workstations. Upgrades are available from 64mb to 512mb and carry a lifetime express-ship replacement warranty.

800-692-4343
Circle No. 42



UNIX laptops from RDI

The newest generation of RDI products features mid-level workstation processors, large high-resolution color displays, removable hard disk storage and a low price premium over comparable desktop workstations. This combination of features enables technical users to easily bring their workstations anywhere, so they're more versatile, more flexible and more productive in their work. It's a new way of working, support and service.

800-692-4343
Circle No. 43

High Data Availability, Superior performance, Fault-Tolerant, Scalable, ZZYZX.

The last word in RAID.

ZZYZX RAID solutions designed and built to meet your needs today and tomorrow

Built-to-order. ZZYX RAID solutions are available for desktop or rack-mount environments. We design systems to meet your requirements and your budget while offering the expandability to grow with the demands of your future data storage needs.

Beginning with our knowledgeable sales staff, **Zzyzx** is committed to working with you to determine the best controller configuration, enclosure size and disk drive capacity needed to create powerful, cost effective RAID solutions that meet *your* specific data storage needs.



From entry-level imaging workstations to high capacity OLTP and video-on demand servers, our level of technical expertise and the scalable nature of our RAID controllers, allows us to provide storage solutions for every application.



ZZYZX PowerRAID features:

- ▶ Ultra SCSI host and Ultra SCSI disk drive channels
- ▶ 8GB-2TB storage capacity
- ▶ RAID levels 0,1,0+1,4 & 5
- ▶ Up to 4 host channels
- ▶ 256MB cache (write-back & read ahead)
- ▶ Superior disk drive media error management
- ▶ Global hot sparing with automatic disk drive rebuilds
- ▶ Up to 32 RAID sets and the ability to run different RAID levels simultaneously
- ▶ Dual controller options with automatic fail-over

Fault tolerance

The assurance of high data availability begins with the use of our reliable RAID controllers, but it doesn't end there. These controllers are installed into highly fault-tolerant cabinets that feature hot-swap/redundant power supplies, hot-swap fans, and hot-swap key-locking disk drives carriers.

Supported environments

Zzyzx PowerRAID subsystems are platform independent and are virtually plug-and-play in any SCSI interfaced computing environment.

ZZYZX service and support

At **Zzyzx**, the level of customer service and technical support is unparalleled. Our commitment is to customer satisfaction and 24x7 systems operation.

So why wait? Find out today how you can get more storage than you need for a whole lot less than you think. **ZZYZX, where more is less!**

Call today TOLL FREE:

800.692.4343

www.zzyzx.com/sunx


5893 Oberlin Drive, Suite: 102, San Diego, CA 92121
▼ Voice: 619.558.7800 ▼ FAX: 619.558.8283 ▼ E-Mail: sales@zzyzx.com

Circle No. 6

WORKSTATIONS
ZZYZX
& PERIPHERALS

Ask Mr. Protocol

by Michael O'Brien



TOM BARRETT

"Security considerations are not discussed in this memo."
– Just about everybody

"Internet Drafts may disappear at any time."
– Rules of the game

"What do you mean we can't claim we're compatible? A standard's a standard, right?"
– Hungry vendor with acid stomach

Your Basic Standard Mr. Protocol

Q: *I have just developed the greatest protocol since the Angel of Death invented Passover. Thousands shall fall before the might of my checksums, and loud shall be the lamentations of those who fail to route my packets. The purity of my header bytes shall bring swift destruction and suffering unparalleled to my enemies. Death! Death and ingrown toenails to the infidels who fail to adopt it! And Victory! Victory unmitigated by sorrow, strife or packet loss to the faithful! ... Er... How do I get it adopted as a standard?*

A: No worries, mate. I'd say that you and the Internet Engineering Task Force are about ready for each other. You've got just the right attitude to engage in fruitful negotiations with these stalwart architects and guardians of the Internet.

Things in the IETF have never been what one would call placid. In a way, this is too bad: There have been more than a few people whose work would have made the Internet a better place,

who have left the Internet standards arena because the interpersonal static in the IETF was too much for them. This has only gotten worse over the years, as the IETF population has increased, and as the IETF meetings have attracted a more diverse crowd.

As Mr. Protocol has noted before, the initial population of the Net included a much higher percentage of actual network implementers than today's. In addition, the network architects, designers and implementers were a more homogeneous population than they are today. Consequently, the path that protocols followed on the journey from inception to wide distribution was somewhat more informal than it is today.

The fundamental instrument for propagation of Internet standards has always been the Request for Comments. RFCs have a fixed format that makes it obvious that the network's bread and butter came from the U.S. government. Every single one has a required header reading "Security." This is there because


the security folks didn't have too much trouble convincing the government and everyone else that network security has to be designed in from the beginning, not slapped on after the fact. That the network's initial design was done by academics is obvious from the fact that in almost every RFC, the sole contents of the "Security" section is the sentence: "Security considerations are not discussed in this memo."

Everyone who wanted to invent a protocol was free to do so. In fact, in large measure, this was what life was all about. Here was this nifty new network; let's invent some services for it! This was simple enough. First of all, you invented the protocol. Then you implemented it, and after you had a working implementation, you fixed it and fussed with it and futzed with it until you had the rough edges worn off, and then you wrote an RFC about it. You also talked about it on the relevant mailing lists and convinced someone else to write an implementation for some other kind of

Is your high-availability network suffering from a low-availability legacy?



Introducing New Symmetra™ Power Array™ from APC.



If you've ever installed a 6-20kVA, non-redundant, legacy UPS system in a datacenter, you know the challenges: installation is a struggle, maintenance is a challenge, power upgrades are unheard of, and hardware tends to be unreliable. In fact, most legacy UPSs were designed for electrical and manufacturing applications and not with the needs of today's sensitive datacenters in mind. Worst of all, with a single point of failure, the legacy system is all too interruptible.

With power problems causing nearly half of all network downtime, the fastest and easiest way to increase network availability is to increase power availability. Now, APC's Symmetra™ Power Array™, a revolutionary power protection system, provides the high availability, flexibility, reliability, ease of installation and simple upkeep you need – all without breaking your back or the bank. Best of all, Power Array redundancy and manageability deliver uninterrupted peace of mind.


N+1 Redundancy: Symmetra Power Modules share the attached load evenly. If one module is removed or fails, the other modules seamlessly and instantaneously pick up the rest of the load. N+1 redundancy is used today in disk arrays (RAID), in processor power supplies, and in processors themselves. Symmetra Power Array brings redundancy to server power protection for the first time.

Scalable Power: Symmetra Power Array can be configured for 4-16kVA, enough for mid-sized computer rooms. Incremental 4kVA Power Modules expand UPS voltage load capacity to match datacenter demands.

Scalable Runtime: Additional Battery Modules, which may be shifted from one Power Array to another, or added in Extended-run Battery Frames, can provide virtually unlimited runtime. If you need to solve the problem of network downtime once and for all, install a Power Array with redundancy and fault-tolerance to match your server and storage. Call APC today for your FREE Enterprise Solutions Kit.

Add Redundancy and Fault-tolerance

True N+1 Redundancy ensures maximum uptime and high systems availability. In a Power Array, battery, power and intelligence modules are all redundant, providing total power reliability for your network.



Add Capacity/Runtime

Protect your investment with redundant, modular power modules: Scalable kVA and runtime ensure that Symmetra will meet your power needs now and in the future.

Simplify Maintenance

Symmetra's service-friendly design cuts the cost and complexity of maintenance. Batteries, intelligence modules and power modules are user-replaceable and hot-swappable.

Integrate Advanced Management

Symmetra power array provides the ability to manage your network or datacenter power from anywhere in the world with PowerChute® plus management software. Whether you use a Web-based or LAN/SNMP-based management platform, APC software and accessories integrate with your existing management solution.



FREE! Enterprise Solutions Kit!

an invaluable set of tools to help plan your datacenter power protection, including the Gartner Group report, "Uninterrupted Power Supply Systems in a Distributed Computing Environment". Just go to

<http://promo.apcc.com> KEY CODE D506Z

or E-mail datacenter@apcc.com and get your FREE kit today!



(888) 289-APCC x7259 • FAX: (401) 788-2797

Installing or reconfiguring your modular Power Array couldn't be any simpler.



©1998 APC. All Trademarks are the property of their owners. SY1A8EF-US

(800)347-FAXX PowerFax

E-mail: datacenter@apcc.com

132 Fairgrounds Road, West Kingston, RI 02892 USA

APC has won more awards for reliability than all other UPS vendors combined.



Circle No. 7

machine. Then when the two implementations refused to work together correctly, you fussed with the implementation and refined the protocol and its description until they did, at which point you declared success and moved on. Except that people wouldn't let you, because like a Chinese person whose life you'd saved, The Protocol was now your responsibility and prospective implementers felt free to wake you up in the middle of the night to talk about pathological cases for the rest of your natural life.

Which is exactly the reward you'd prayed for, right?

The only tricky part here was the protocol number. In ARPANET days, you needed a socket number. For the Internet, you need a port number. Same thing, different name. It's a way of identifying just what user-level protocol this

TCP packet is supposed to be. Every protocol on the Internet, from highest to lowest, has some sort of identifying number that makes it unique among protocols at that layer. User-level protocols like FTP, Telnet and the Web have port numbers assigned to them. Telnet is 21, FTP is 23 and the Web is 80. You can prove this with Telnet. Telnet to another machine's port 23 and you'll get an FTP prompt. Telnet to port 25 and you'll get an SMTP prompt (25 is the mail port). Telnet to port 80 and you'd better be ready to type some HTTP stuff. Below the user level, TCP and the Internet Control Message Protocol (ICMP), each have identifying numbers. IP itself is identified by a number, including a version number (currently 4, hence all the talk about "IPv4" and "IPv6," which is coming along to replace it).

The Internet Assigned Numbers Authority (IANA) assigns protocol numbers. In the early days, getting a new number was a simple matter of sending an email saying, roughly, "Gimme a number." You had to provide some information, such as the title of the protocol, your name and contact

information, and a note from your mother saying that yes, this was a genuine protocol that people were actually going to use and not a master's thesis with a case of gigantism. And that was pretty much it. If IANA didn't hear from you for a while, you might eventually get a note saying, "Is this thing still real, or can we release the number now?" Aside from that, you pretty much owned the road as far as that particular protocol number went.

On the ARPANET, congestion was the only real problem the ordinary network user, or even the ordinary protocol implementer, could cause.

Now, it might be thought that if this were a network upon which people depended, allowing such free and easy experimentation might not be such a good idea. For a number of reasons, this cavalier attitude was not only allowed, but encouraged.

For one thing, that was the purpose of the net. The entire network was experimental, even though people depended on it. Eventually, the MILNET was split from the ARPANET, a year or two after the TCP/IP protocol suite replaced the NCP protocol, thereby allowing multiple networks with routing between them. It was this event that marked the first retreat from a regime of more or less untrammelled experimentation.

Second, at least on the ARPANET, it was hard to break the net. Practical knowledge of packet-routing algorithms was still a dicey proposition, but it wasn't a proposition that users ever encountered directly. The ARPANET was a two-level net. Only the Interface Message Processors (IMPs) communicated directly over the net, and they were homogeneous, centrally maintained and shared a common base of routing software. Users didn't have it in their power to alter or affect this level of the network in any way, aside from an unfortunate proof-of-concept escapade at Massachusetts Institute of Technology that partitioned the net for a day or two. Congestion was the only real problem the ordinary network user, or even the ordinary protocol implementer, could cause.

Not that they didn't try. Before the

IP protocol was deployed, institutions could connect only one or two mainframes to the net, that being all that the IMPs could support. As the brand-new concept of local-area networking spread through the community, more than one institution seriously proposed various sorts of gimmickry to allow the host connected to the IMP to act as a gateway. NCP couldn't support this in any straightforward manner, there being no such thing as a network number in NCP. The schemes involved doing various sorts of violence to the host address. These ideas were rejected. As the users didn't actually have any access to the internals of the IMPs, that was the end of that idea.

Enter Internet Standards

As the Internet has grown, the "standards path" has become more elaborate. Initially, there wasn't one. Everything was an RFC, plain and simple. Some RFCs (the one specifying IP, for example) were more important than others (the one for obtaining time and temperature in the machine room). The fact that there were no standards was regarded by some as a strength, because it allowed people to try alternative solutions to established practice without sanctioning hordes of people crying out that "standards are being violated."

As more and more people's livelihoods came to depend on the Internet, and to make it easier to deal with international bodies—which believe that standards are fundamental—the Internet community eventually began to evolve its own standards track process. The RFC describing the process has undergone two revisions since it was first issued.


Some things remain the same. First of all, RFCs are not synonymous with standards documents. Not all RFCs are standards documents. Anybody can still write an RFC and submit it for publication. However, very few RFCs are now submitted out of the clear blue sky. A new type of document, the Internet Draft, has been created to allow work in progress to be published.

Internet Drafts are peculiar beasts. If left alone, they disappear automatically and tracelessly. Consider this a type of intellectual garbage collection. They

full **SUN** compatibility

THOUSANDS installed

TATUNG SBus or PCI



SBUS


With 10's of thousands of high quality workstations and servers installed in North America, Tatung is second only to Sun in SPARC®-based computing. Tatung U-series workstations feature 100% Sun UltraSPARC™ microprocessors in single or dual processor configurations and your choice of SBus or PCI versions. All systems are fully compatible to Sun and are priced for less.

Tatung was licensed to manufacture Sun compatible workstation and servers more

than 10 years ago. This commitment to SPARC-based computing has led them to an impressive line of UltraSPARC and microSPARC™ systems. All Tatung solutions include the highly-respected multitasking, multithreaded Solaris® OS at no additional cost.



As with all Tatung workstations, the SBus and PCI UltraSPARC U-Series are manufactured in the US. Call us today to power up a system for you. 1-800-659-5902



PCI

Contact Tatung for Internet/Intranet Server Solutions at 1-800-659-5902.



Tatung's full line of UltraSPARC and microSPARC workstations and servers

©1998 Tatung Science and Technology, Inc.

TATUNG[®]
SCIENCE & TECHNOLOGY, INC.

The Intelligent Choice in SPARC Computing Solutions

Phone 1-800-659-5902 • Fax (408) 383-0886

(408) 383-0988

email — mkt@tsti.com or website — www.tsti.com

All trademarks are property of their respective owners.

18GDA

Circle No. 8

may be revised at any time, replaced at any time or pulled at any time. The “sense of the Senate” is that they may not be quoted in any paper, publication or report, and they may never be referred to in the sense of a standard to be complied with. The aim here is to prevent the Internet implementers from inventing some new protocol, and while the protocol is still in the fiddling and futzing stage, have some vendor suddenly and loudly announce that they support it. Of course this happens anyway, but in general, vendors don’t refer to Internet Drafts directly, otherwise their own representatives in the standards process will be given a hard time (which they are pretty much guaranteed to pass along within their company).



The normal way in which a standard comes into existence these days is that a working group is formed within one of the “Areas” of the IETF. The IETF meets three times a year, and working groups commonly arise out of “Birds of a Feather” sessions held at these gatherings. If the BoF session decides that a working group is warranted, they work with the area director to form one.

Once formed, the working group creates a mailing list and begins work on defining the protocol specification. Working groups, although the fundamental bodies for the creation of new standards, are themselves nonstandard. Internet architects and implementers include in their number a fair proportion of colorful individuals, and people skills are not high on the requirements list for membership. If physical assaults are rare, it’s only because these people see each other a maximum of three times a year.

How does any work get done? Mr. Protocol is glad you asked.

The mailing list, alternating with IETF meetings, generally gets the job done despite the occasional noise. Once in a while, a working group fails due to internal politics, and the ensuing stench is great. This is a remarkably rare event, all things considered. Generally, people either get tired of arguing, or get dis-

gusted and leave, or (mirabile dictu) work out their differences until an Internet Draft is created specifying the protocol.

At this stage, the working group chair requests that the Internet Draft be entered onto the standards track as a Proposed Standard. The IETF proper doesn’t have much of a structure aside from the area directors, so oversight is given to the Internet Engineering Steering Group, a much smaller body elected from the membership. The IESG oversees all advances along the standards track, and if it approves, it instructs the RFC editor to advance the draft to an actual RFC.

A Proposed Standard has to be pretty much complete. In particular, it’s a big boost if it’s actually been shown to work in the real world, without breaking the world in

consequence. Once two independent implementations have been created and real-world hang-ups fixed, the IESG may order the specification to be advanced to Draft Standard. At this point, the main thing standing between the spec and life as a full Internet standard is a large amount of real-world experience, together with a show of popularity. At each stage of the specification’s progress, a new RFC is issued. RFCs, once issued, never change. That way, RFC 2036, for example, is always going to be the same. It may not be current, but it will always have the same text.

How to tell what’s current? By giving RFCs two names. There is a parallel collection of STD documents, which represent the current standards. The contents of these do change over time. The change is made by pointing the current STD label away from the old RFC and toward the new RFC. The resulting document has a new RFC number but the same old STD number. There is an STD document, periodically updated, which gives the current status of all documents on the standards track.

The third version of the RFC detailing the standards track makes interesting reading. A sizable chunk of it deals with a staged process of appeals, going up through the IESG and the Internet

Activities Board all the way to the Trustees of the Internet Society. None of it talks about lawsuits, which is too bad, because if the bozos currently attempting to litigate their way to fame and fortune by suing the IANA were thereby to lose their standing in the standards process, life might be a great deal simpler...at least for the IANA (and probably, eventually, for the rest of us).

Standards don’t last forever, and those whose usefulness has passed are retired to Historic status. As the standards track RFC points out, the proper word to use here would be Historical, but the use of the word “Historic” is now itself historical.

Room to Experiment

Finally, there are protocol specifications that are not on the standards track, will never be on the standards track and which basically should never come near the standards track. The bin for these things is labeled “Experimental,” and if that word strikes terror into your heart, you’ve demonstrated good perception. This is where most of Mr. Protocol’s own new protocols wind up. The protocols in this category are either intended for a narrow use, are driven by some other standards body or are just plain outright dangerous.

Consider a superset of TCP, designed for use with satellite communications. Now, ordinary TCP, although it doesn’t say so in the spec, backs off and slows down whenever it detects packet loss because it assumes that in today’s Internet, substantially all packet loss is due to congestion. Well, in the satellite world, most packet loss is just due to loss: When one bit in 10^6 or so goes flooey, you’re going to get dropped packets and that’s that. The correct response in this case is not to back off. If anything, it’s to speed up for a while. Detection of the dropped packet requires that the network pipeline to the satellite be drained and refilled, so the faster the source can do that, the better. In fact, an experimental protocol for satellite transmission has been created that does exactly this but still behaves well in genuinely congested scenarios. For various reasons, this protocol has a wardrobe of different

The PC X Server
Standard:

Exceed



www.hummingbird.com/sem (416) 496-2200

Recent Excellence Awards:



The Gold Standard:

Exceed Version 6

The world leader now brings the world to you. Exceed, the fastest growing PC X server, outsells all other PC X servers combined, with over 55% of the world market. We packed the new Exceed Version 6 with highly enhanced features:



new web capabilities, improved performance, advanced functionality and an enriched and simplified interface. Then we added HostExplorer Terminal Emulation Suite including TN3270E for today's enterprise computing environment. Exceed Version 6 for Windows NT and Windows 95 and Exceed Version 5.2 for Windows are all on one CD. Being the best-selling, most powerful X server on earth just wasn't enough. It's only natural we keep exceeding... it's in our name.

discover everything net

to build your business

Featuring:

E-COMMERCE SYMPOSIUM
INTRANET FORUM
NETWORK SECURITY DAY
KNOWLEDGE MANAGEMENT FORUM
INTERNET PROMOTIONS WORKSHOP
ISP WORLD
TCP/IP WORKSHOPS
PLUS, 100 OTHER SESSIONS!

Special Focus:

WEB DESIGN & DEVELOPMENT SERIES
WEB DESIGN FORUM
XML XPOSED
WEB PUBLISHING DAY
JAVA DEVELOPER DAY
E-BUSINESS WITH VISUAL BASIC
E-BUSINESS WITH LOTUS/IBM
DESIGNING EFFECTIVE WEB SITES
WEB DATABASE WORKSHOP

CONFERENCE INFORMATION:

Visit: www.internet.com
Call: 1-800-500-1959
E-mail: iwsprogram@mecklermedia.com

TO EXHIBIT:

Call: 203-341-2923

THE SECOND ANNUAL

SUMMER

Internet



World 98

CONFERENCE & EXHIBITION
JULY 13-17, 1998
McCORMICK PLACE
CHICAGO, ILLINOIS



PRODUCED BY:

Mecklermedia
THE INTERNET MEDIA COMPANY

CO-SPONSORED BY:

Chicago Tribune
chicago.tribune.com

SPONSORED BY:

**internet
WORLD™**

internet.com

NetSearcher™
<http://www.netsearcher.com>

**WEB
DEVELOPER.COM**

SHOW NETWORK
PROVIDED BY:

IBM

From the E-mail Center to the Show Floor, IBM puts the "Net" into Internet World.

Code: SEJ

Ask Mr. Protocol

faces it can put on, only one of which looks like TCP. Given all this, and given the fact that it's not designed to replace TCP globally, now or ever, you have a perfect candidate for "experimental" status.

Perhaps the most interesting feature of the IETF standards process, aside from the fact that it is continuously evolving as people discover new mud pits to fall into, is the degree to which it has managed to preserve the network's egalitarian beginnings. Although most of it is taken up with the humdrum stuff of working group creation and management of the group's output, it also documents a second path along the standards track for standards that are not the output of a working group.

These individual contributions have proven highly useful in the past, so a path remains for the individual to propose a new protocol, independent of the IETF process. The RFC editor may suggest that the individual share the work with a working group if that seems appropriate, but if no appropriate working group exists, and the individual cannot or does not want to pull together a working group, the document can move along the standards track like any other, as long as the IESG agrees to promote it.


It gives Mr. Protocol a warm and happy feeling to know that his own personal efforts still have a standards track outlet. →

Mike O'Brien has been noodling around the UNIX world for far too long a time. He knows he started out with UNIX Research Version 5 (not System V, he hastens to point out), but forgets the year. He thinks it was around 1975 or so.

He founded and ran the first nationwide UNIX Users Group Software Distribution Center. He worked at Rand during the glory days of the Rand editor and the MH mail system, helped build CSNET (first at Rand and later at BBN Labs Inc.) and is now working at an aerospace research corporation.

Mr. Protocol refuses to divulge his qualifications and may, in fact, have none whatsoever. His email address is amp@cpj.com.

Who Is Responsible for Thousands of UNIX Solutions on Windows NT ?



NuTCRACKER

(that's who!)

Already, NuTCRACKER® has produced thousands of UNIX® software solutions on Windows NT® for customers in Federal Government, manufacturing, oil & gas, CAD/CAM, telecommunications, retail, financial, software and other industries. The reason? NuTCRACKER lets UNIX programs run natively in the Windows® environment.


NuTCRACKER extends the powerful Win32® subsystem, allowing rich integration between UNIX and Windows programs. That makes it the perfect solution for Windows 95/98 and Windows NT users. Plus, NuTCRACKER's "write once—deploy many" capability supports full, cross-platform development. When your source program is written once on any brand of UNIX—it can run on any Windows NT/95/98 workstation or server.

Developers love how much NuTCRACKER cuts development time—70%, 80%, even 90%. And there are thousands of customers and case studies to prove it. Call today and see why NuTCRACKER should be your UNIX to NT solution.

To find out about our
30-day risk-free trial offer, call today.

800-637-8034 toll free
703-803-3343 Washington DC Metro

visit us on the web at: www.datafocus.com



The Enterprise UNIX Plug-in for Windows

NuTCRACKER is a registered trademark of DataFocus, Inc. Windows, Win32, Windows 95/98 and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the U.S. and/or other countries. All brands or product names are or may be trademarks of, and are used to identify products and services of, their respective owners.

UNIX Basics

by Peter Collinson, Hillside Systems



DENISE ORTAKALES

make: Mastermind of the Update

It was interesting to read Rich Morin's I/O opener article "Page Processing in Perl" in *SunExpert* (March 1998, Page 43). It turns out that I have reached a similar conclusion about generating Web pages. Like Rich, I've also written a processor routine in Perl that takes page descriptions and spits out pages.

Perhaps four years back, I realized that each of my Web pages has a large chunk of HTML that remains essentially the same across any particular set of pages. The constant part usually provides a consistent look and feel across a whole Web site, or part of one. I now create pages by writing a small Perl program that fills in a template taken from a file, where the variable sections in the template are supplied from the Perl code. Creating these little Perl program description files is easy, because all you supply is the information that differs on each page. Also, if I want to make global changes to the site, I can alter the template file and recompile the pages. I can change the site without the massive

time-consuming, error-prone task of editing each page individually.

My processor is designed to be a simple classical macro processor, although it's not that elegant and only does a bare-bones job. The processor scans the source file for a name (enclosed in double braces, `{{CONTENT}}`, for example) and replaces any embraced name it finds with contents of a variable in the Perl program (taken from an associative array, for Perl-literate readers). There are also simple macro definitions and conditional facilities.

I seem to be using this beast on every Web site I design. It can compile static pages from simple short Perl programs that are applied to a template file, provide dynamic pages from CGI scripts in response to input from forms and add addressee information and other material into stock email messages.

However, the point of this article is not to discuss the ins and outs of converting Perl programs into HTML documents, but to look at the way I've


automated the compilation system that uses it. I want to ensure that if I change a page description, I'll run the appropriate programs on the appropriate files to generate the updated pages. The update task is masterminded by a standard UNIX tool: `make`.

I've discussed `make` in this magazine before. I wrote two consecutive articles, "Make: Parts I and II," February 1992, Page 34, and March 1992, Page 26. This was back in the dim, distant past for many readers, so I make no apology for revisiting some of the details of `make` again—although this article does contain information that was not discussed previously.

make: The Basics

The notion that underlies `make` is easy to understand. Many files on a UNIX system are created by applying a program to a set of source files. It's a simple observation that the modification time stored with each generated file must be later than the time stored for

Is your data (like the Black Rhino) an endangered species?



Save your data with Hitachi Freedom 5700E and 7700E Storage Solutions.

Every single day, your company's data is threatened with extinction. Which is why you should consider all the ways that Hitachi Freedom Storage™ can save you:

- Industrial-strength RAID for ultra-high-speed performance
- Continuous data access across multiple platforms
- Centralized management software for enterprise-wide control from a single location
- Customizable storage that's flexible and scalable
- Unique, integrated hardware and software based solutions for Year 2000 testing and validation



Visit our Hitachi Data Preserve and we could send you to a Wildlife Preserve. Call us today or visit us on the Web to request a no-obligation assessment of your enterprise storage needs, then register for our "Safari West" Adventure Giveaway.

1-888-333-4488, ext 111, or www.hds.com/storage/datapreserve.html

HITACHI
DATA SYSTEMS

No purchase necessary. Sweepstakes open only to legal residents of the United States. Void in the States of Florida and Rhode Island, and in any other jurisdiction where prohibited by law. The winner will receive a 7 day trip for two to Sonoma County, California, including round trip airfare and a visit to the "Safari West" wildlife preserve. The estimated value of this trip is \$3000.00. Entries must be received by July 31, 1998. Odds of winning depend on number of entries received. All federal, state and local laws apply. Hitachi reserves the right to modify, cancel or suspend these sweepstakes.

each of its component source files. In make-speak, each generated file *depends* on the set of source files.

If we discover that the modification time on one of the source files is later than that on a generated file, then the generated file is out-of-date and needs to be updated. The make command is fed a user-generated control file that provides a specification of how a particular generated file, the *target*, depends on a set of source files. The control file contains a *rule*, or a set of rules, used to convert each source file into a target file. Unless it is given a specific file name, the make command first looks for a control file called `makefile`, then for a file called `Makefile`. You can choose which flavor you prefer. Most people these days use the capitalized name, and many versions of make now complain if both files are present in the same directory. To avoid confusion, I'll call one of these files a "makefile."

The makefile will contain several definitions of the form

```
target: dependency list
<tab> rule
```

The syntax tells make that the particular *target* depends on a list of files, and to create the target the *rule* should be executed. In most versions of make there must be a tab character before the rule, and I've emphasized that here, I'll be showing white space in the examples that follow. The rule is simply a shell command that is run by make. It is assumed to create the target. In general, if the command fails returning an error status, then make will spot this and terminate any further execution.

Actually, the rule section can contain several commands, appearing on several lines. Each command is run in a separate shell, and this can catch you out sometimes.

Let's look at a specific example. The easiest way to show you what happens is to pick up a programming task, but the strategy can be applied to many nonprogramming jobs, so don't think this article is aimed just at programmers.

If we are compiling a C program from a source file called `hello.c`, then we might create a makefile that contains

```
hello: hello.c
    cc -o hello -O hello.c
```

The target is a compiled program called `hello` that depends on the source file `hello.c`. If the modification time on `hello.c` is later than that on `hello`, then we will invoke the rule. The C compiler is called and told to place its output into a file called `hello` (the `-o` option). We've also asked for the compiler to run the optimizer as part of its work (the `-O` option).

I should emphasize that if the target file `hello` was created more recently than `hello.c`, then the rule is not run by make. However, if we omit the dependency, starting the line with

```
hello:
```

then the rule will always be run. The trick with constructing

makefiles is to get the dependencies right.

The example above is a no-brainer. There is no huge win when compiling a single file, except that you can type make as a reflex action and not have to worry about how the program is compiled. Actually, the ability to just type make can be a big win, saving a lot of time in understanding just how the command is compiled.

However, the real power of the make command is apparent when there are several source files to be compiled into the final target. For example,

```
hello: a.o b.o
    cc -o hello a.o b.o

a.o: a.c header.h
    cc -c -O a.c

b.o: b.c header.h
    cc -c -O b.c
```

The file starts with the final target, the `hello` program that depends on two compiled modules, `a.o` and `b.o`. We place the final target first because by default make will create the first target that it finds in the makefile. So when we type make, a tree of dependencies will be built, and that tree is then traversed looking for work to be done.

If the targets `a.o` and `b.o` are not current, they are each created from their own rules. This will result in separate runs of the C compiler that is given the `-c` switch telling it to create a compiled module. Notice that each compilation depends on the `header.h` file. If I change the header file, then both `a.c` and `b.c` will be recompiled. When `a.o` and `b.o` exist and are more recent than the current version of the `hello` file, then the rule to create the final target will be executed.

Incidentally, I mentioned that make could be given targets to work on, so we can type

```
$ make a.o
```

to run the system to create that particular intermediate step. This feature is perhaps more useful if you have several programs to be compiled in the same makefile:

```
all: prog1 prog2

prog1: prog1.o proghdr.h
    cc -c -O prog1.c

prog2: prog2.o proghdr.h
    cc -c -O prog2.c

and so on
```

Now we can type make to create both programs and select one or the other by supplying a specific request on the command line. This makefile also demonstrates that the rule section can be empty.

In many makefiles, you'll find a bunch of ad-hoc targets

RELIABLE OVERLAND DLT™ AUTOMATION

Backup and restore your peace of mind.

Relax. It's Overland. Our full line of DLT tape solutions is designed to make your job easier and more productive. Emphasis on reliability. Industry-leading warranties. Commitment to high performance. Compatibility with all leading backup software. Focus on the right price. Scalability to make your investment last. That is what customers around the world have come to expect from us. And there's more to come. The seven major industry awards for excellence that we have won are just the beginning.



The LXS MiniLibrary™ is cost effective DLT automation for workgroup to enterprise data storage. Recommended for systems with 50 GB to 350 GB of data. Available in either desktop or rackmount. Starting at \$10,495.



The DLT LibraryXpress™ offers true scalability. It grows as you grow, offering more capacity and performance options than any other library. For systems from 50 GB to 3,000 GB. Starting at \$13,995.



The New LoaderXpress™ provides DLT automation - at a cost comparable to a stand-alone drive. And a 5- or 10- cartridge magazine allows for up to two months of unattended backup. For systems from 15 GB to 100 GB. Starting at \$4,995.



Relax. It's Overland.

www.overlanddata.com
1-800-729-8725 or 619-571-5555

Worldwide 619-571-5555 Europe (+44)118-9891891

DLT is a trademark of Quantum. All others are trademarks of Overland.

Circle No. 13

that perform useful functions:

```
clean:
    -rm *.o core hello
```

The `clean` target is used to ensure that created files in the directory are removed. The hyphen preceding the `rm` command is a small bit of magic. Normally `rm` will complain and return a failure status when it is given a file to delete that doesn't exist. The nonzero status will cause `make` to terminate because one of its children failed. We can tell `make` that we can tolerate failure by placing a hyphen before the command. Quite often, I'll add the `-f` flag to an `rm` command in a makefile because the `-f` flag changes the default behavior; the command doesn't complain and die if a file doesn't exist.

Other common targets in makefiles are: `install` to place the final binary into public use in the file system; `dist` to create a distribution, perhaps using the `tar` command to create a file that can be transported elsewhere; `print` to print the sources; and `depend` to create a dependency set that is appended to the makefile.

More Syntax

Well, all the above concentrates somewhat on programming, but we can use `make` for any purpose where the modification time of files is important. For example, let's say we are working on several important files in a directory and want to create a checkpoint copy from time to time. We create a subdirectory called `checkpoint` and write a makefile like this:

```
FILES= file1 file2 file3 file4 file5

check: $(FILES)
    cp $? checkpoint
    touch check
```

We type `make` and only the files that we have altered since the last run will be copied into the `checkpoint` directory. I've introduced some new things into this makefile. First, you can define macro replacements by using the `NAME=list` statement. When `make` encounters a `$(NAME)` elsewhere in the makefile, the text is replaced by the contents of the list. So in the example above, our main target, `check`, depends on all the files in the list `FILES`.

The dollar syntax is used to introduce other results of the dependency check. Here we are using `$?` , which is replaced by the list of files that are newer than the target. There are other magic dollar values that I'll get to later. So if `file2` and `file4` are newer than the target, and we type `make`, then the `cp` command becomes

```
cp file2 file4 checkpoint
```

The target is a file called `check` that is there simply to record the time that the last checkpoint copy was made. It's created by the `touch` command, which exists on all UNIX systems and simply changes a named file, updating its modification time.

We should probably add a way of making a complete new checkpoint copy to the makefile instantly:

```
checkpoint:
    cp $(FILES) checkpoint
    touch check
```

so we can now say

```
$ make checkpoint
```

and make a complete copy of all the files. Also, in this type of makefile, I like to add something that I can use to obtain all the names of the "interesting" files:

```
names:
    echo $(FILES)
```

which means I can use the backquote operator in the shell using the names. The command

```
$ grep fred `make names`
```

will look for `fred` in the files that are relevant in the current directory.

Implicit Rules

To make things somewhat easier for programmers, `make` comes with a set of default rules that are generally loaded when the command starts. The default rule set used to be part of the binary of the program, but in recent years it has migrated to a file stored in some public place; it's `/usr/share/lib/make/make.rules` on my Sun running Solaris 2.6.


The default rules work with file suffixes, so the rules understand that a file ending in `.c` is a C source file, a file ending in `.o` is an object file containing a compiled module and so on. The `make` program itself knows only about suffixes and is given a strong hint of the order in which files will appear by the contents of the special target `.SUFFIXES`. The target is given all the suffixes that the rules know about. A simplified default version for C programmers might be

```
.SUFFIXES= .o .c .h
```

The order of the list is important. Target suffixes come first, so programs are made from `.o` files that are created from `.c` files. C program files can also include header files that will end in `.h`. The default file will contain a set of special rules that specify how to turn a file with one suffix into a file with another, so to make a `.c` file into a `.o` file you will see a rule like this:

```
.c.o:
    $(CC) $(CFLAGS) $(CPPFLAGS) -c $<
```

The various macro definitions here will be established with default settings, so `CC` will be `cc`, and `CFLAGS` and `CPPFLAGS`



There's only one thing spreading more quickly than Year 2000 madness, and that's a proven solution to the dilemma - EMC Enterprise Storage.™ With EMC TimeFinder™ you can quickly and easily segregate

It will have you reaching

test data from production data, without monopolizing host resources, thus reducing the time to Year 2000 compliance. And insuring continuous, uninterrupted business operations throughout conversion,

for the champagne,

testing and migration. In addition, EMC offers the first storage architecture to simultaneously support mainframe, open systems and NT environments. So once you've achieved Year 2000 compliance, your

instead of the aspirin.

company can redeploy its EMC solution for maximum return on investment. The next millennium should give you reason to celebrate, not hyperventilate. Call 1-800-424-EMC2, ext. 302. Or visit www.EMC.com.

Circle No. 14

EMC² is a registered trademark and EMC, EMC Enterprise Storage, The Enterprise Storage Company, The EMC Effect and EMC TimeFinder are trademarks of EMC Corporation. ©1998 EMC Corporation. All rights reserved.

EMC²
The Enterprise Storage Company

are set to null strings. The purpose of these definitions is to permit the user to set values in their own makefiles that establish local definitions for the default rules. The `$(` at the end of the line is another magic variable that is replaced with the name of the dependency file. This rule can now be used whenever it's necessary to create a `.o` file from a `.c` file, so our makefile to create `hello` from `a.c` and `b.c` can be made much simpler:

```
OBJS=a.o b.o
CFLAGS=-O

hello: $(OBJS)
    cc -o hello $(OBJS)
```

Notice that there is no explicit mention of the `.c` files. We've told `make` that it needs two object files and it will use the suffix rules to look for candidate source files. When using the full default set, these candidate files can be written in C (with a `.c` suffix), in C++ (with a `.cc` suffix), in FORTRAN (with a `.f` suffix) and so on. In this case, we are programming in C, so `make` will successfully find `a.c` and `b.c`, apply the compilation rule (using our version of `CFLAGS`) and create the program.

Making Your Own Rules

Well, I started this article by discussing how I created HTML files from a Perl source, so let's return to that example. We can, of course, create makefiles that contain specific rules telling `make` how to create a `.html` file from a `.pl` file. This looks something like

```
page.html: page.pl
    $(PERL) -w page.pl > page.html
```

where we set the `PERL` macro value to be the location of the Perl binary on the system. The temptation is to cut and paste this first definition whenever we add a new page into the makefile and change the bits that need editing to specify how to create the new page from the brand-new Perl source file. However, in the long run, it's much simpler to create a rule and use the ability of `make` to deduce things about the files it needs from the rule. The makefile looks like this:

```
TARGETS=m.html p.html

PERL=/usr/local/bin/perl

.SUFFIXES: .html .pl

.pl.html:
    $(PERL) -w $(< > @$@

all: $(TARGETS)
```

The first line sets up the names of the HTML files that we

are creating. When we want to add a new file, all we do is add a new target name to the list. In fact, the makefile only has one target, `all`, and we tell `make` that it's dependent on the files we need to create.

We use the implicit rule that is defined in the middle of the file to do the work of creating the target files. We must first tell `make` that we will be making a `.html` file from a `.pl` file, and we do this by setting two new suffixes into the `.SUFFIXES` target. Actually, this adds the two new suffixes to the current extant list, so to clear the list we'd need to say:

```
.SUFFIXES:
.SUFFIXES: .html .pl
```

Once we have defined the two new suffixes we're using, we can define a rule that is used to create a `.html` file from a `.pl` file. This runs the Perl interpreter with the `-w` option to help with error checking—it takes its program from a file in the dependency list (`$(`) to its equivalent target (`$(`).

We are now in a situation where any change in a `.pl` file will result in the appropriate `.html` file being rebuilt when `make` is typed.

However, I've also said that the Perl script uses

a template file, and it would be a good idea to rebuild the pages whenever that template file is altered. We need to add an explicit dependency to force a rebuild. We can do this by adding

```
$(TARGETS): template
```

to the end of the makefile. Now when the template file changes, all the files in the `TARGETS` list automatically become out-of-date and are rebuilt.

The good thing about creating a rule that applies to known suffixes is that the makefile becomes very simple. We can add new target files by adding their name to the `TARGETS` list. The makefile can be easily reused in new parts of the system by copying it and replacing the `TARGETS` list with the new targets that are appropriate for that directory.

Finally

If you are running Solaris, and type `make`, and get a "Command not found" message, then you need to be aware that the `make` command lives in `/usr/ccs/bin` by default. You need to change your search path to include that directory. Check that the directory exists on your system first, and if not, politely ask your systems administrator to load it for you. It's an option to the standard installed system. ➔

Peter Collinson runs his own UNIX consultancy, dedicated to earning enough money to allow him to pursue his own interests: doing whatever, whenever, wherever... He writes, teaches, consults and programs using Solaris running on a SPARCstation 2. Email: pc@cpq.com.

The good thing about creating a rule that applies to known suffixes is that the makefile becomes very simple. We can add new target files by adding their name to the TARGETS list.

Running Into Potholes...



On Your Road to NT and UNIX Connectivity?

Running into trouble trying to integrate your UNIX and Windows NT network? TotalNET Advanced Server (TAS) software will smooth the rough road ahead.

TAS enables UNIX computers to become NT file, print and application servers. Setting up and using TAS is quick and easy, thanks to intuitive, browser-based installation and graphical configuration wizards. No additional software is necessary on the NT workstation!


NT File/Print/Application Services: NT users access files and printers residing on UNIX servers using normal NT functions. TAS also enables NT users to access NT applications stored on a UNIX server.

Common File System: Data and applications are stored in a central TAS-based server where NT and UNIX users can easily access the same data.

Transparent to the NT Desktop: The TAS server is seen by NT users as a PC server, so users do not need to know UNIX to access resources on that server. No retraining is involved.

Scalability: TAS provides file/print/application services to thousands of NT and PC workstations.

Is your current solution steering you down the wrong road? Upgrade to TAS and save substantial time and money. Call today for details!



© 1998 Syntax, Inc. All Rights Reserved.
All trademarks of companies whose products are referenced are hereby recognized.

SYNTAX

2 5 3 8 3 8 2 6 2 6
<http://www.syntax.com>



A Special Kind of Freeware

I've been using freely redistributable software for most of my 30-year career in computer programming. I've found DECUS and USENIX tapes to be handy and I've also played with Icon, Perl, Software Tools (RatFor) and UCSD Pascal, among others.

During the last 15 years or so, my involvement has been more active: I edited a couple of Sun User Group (now defunct due to user apathy and Sun's opposition) tapes, then started Prime Time Freeware, a publisher of freely redistributable software.

Through the years, I've had to wrestle with many naming problems involved with "free([ly redistributable] soft)ware." I have answered email explaining why I was not required to send free CD-ROMs to random individuals, written Web pages and columns on the topic, and spent time talking to Richard M. Stallman, founder of the Free Software Foundation, about whether I should use the term "freeware," "free software" or "freely redistributable software." I agree with Stallman that the word "free" is a valuable reminder of the

freedoms this kind of software protects. If software is freely distributed in source form, I'm free to examine it, modify it and, in general, redistribute it to others. I am also free from worry that a vendor will go away (or simply decide to abandon a product), leaving me with an "orphan" piece of software.

On the other hand, "freely redistributable software" is too big a mouthful for many situations. Moreover, the terms "freeware" and, to a lesser extent, "free software" have been adopted by the Macintosh and PC communities for programs that are freely distributed, but only in binary form.

As you can imagine, this naming problem goes right to the heart of Prime Time Freeware's "corporate image." Lacking a better alternative, however, I have simply had to wrestle with the many explanations, none simple or easy.

Flavors of Freeware

Making matters worse, I have found myself dealing with many different "flavors" of freeware. Some authors simply

toss their source code out on a convenient USENET group, (possibly) giving their name and email address as a point of reference. Because of international copyright law, the author retains a real, if limited, copyright to this material. So, if anyone wants to redistribute the code, he must first contact the author for permission (sigh).

Almost as troublesome are the tossed-together notices that some hackers affix to their code: "This code is in the public domain, but you can't charge money for it or use it to operate missile silos." Right.

True public domain (PD) software isn't a problem for users or redistributors. It contains an explicit notice placing the code in the public domain. This allows any recipient to use it, modify it, redistribute it and so on. On the other hand, PD software can be modified, then "claimed" by someone as their own. Worse still, it can be used as the basis for a proprietary package, taking the revisions out of the public domain. Many universities, including Carnegie Mellon, Massachusetts Institute of Technology and

University of California at Berkeley, deal with the first problem by requiring that their code retain a copyright notice. They do, however, allow modification and even proprietary use of the code.

The second problem—proprietary use—led Stallman to craft the GNU's General Public License (GPL). The GPL "protects" the software from proprietary use by requiring that any binary distribution be accompanied by an equally available source-code distribution and that no further restrictions be placed on the modified version. Although you might think the differences in these licensing tactics are pretty minor, they matter greatly to the folks involved. Some free software developers avoid using GNU code, lest portions of their releases become "contaminated" by the GPL and, thus, unusable in proprietary products.

If freeware were used only by occasional hackers, none of this would matter much to corporate types, let alone to the general public. In reality, however, large parts of the Internet rely heavily on free software. For example, the Domain Name System (DNS) is dominated by the Bind package; Internet email forwarding is dominated by Sendmail; Apache runs most Web servers; Perl, Python and Tcl/Tk are very popular languages around the Internet; and C and C++ (via the GNU C Compiler and the GNU debugger) are also part of the freeware opus.

More to the point, the 4.4BSD court proceedings revealed that something like 50% of UNIX System V is based on code developed by research laboratories and universities. This code is now available as 4.4BSD-Lite and many derivative releases.

Netscape Communications Corp.'s recent decision to free the source code for Communicator has, however, brought the media's eyes to the free software community. Freeware has become "hot" and may well get hotter. As a result, the naming problems and internecine squabbles of old are becoming more expensive.

Tim O'Reilly's Summit


Tim O'Reilly of O'Reilly & Associates Inc. is keenly aware of these issues. O'Reilly & Associates publishes a large number of books on free software and related topics. The company also spon-

sors an annual Perl conference and pays Larry Wall to keep Perl alive and growing. In short, Tim O'Reilly is strongly motivated to see the free software community prosper.

Tim recently decided to hold a small "summit" to discuss free software issues. Most of the attendees were developers and included Eric Allman (Sendmail), Brian Behlendorf (Apache), John Ousterhout (Tcl/Tk), Tom

Paquin (Mozilla), Linus Torvalds (Linux), Guido van Rossum (Python), Paul Vixie (Bind), Larry Wall (Perl), Jamie Zawinski (Mozilla) and Phil Zimmerman (PGP). A few others, connected to these developers through business relationships, also attended. To add some differing perspectives, Tim also invited John Gilmore (founder of Cygnus Solutions), Eric Raymond (author of "The Cathedral and the

Data Backup & Recovery From Workstation Solutions



No more endless searching. No more playing hide and seek with your data resources. Just the easiest and most efficient way to backup and recover data – period. We can't say it any more clearly. Quick Restore from Workstation Solutions offers an extremely intuitive, easy-to-use GUI that lets you find backed up data in seconds as well as schedule and monitor backups with the simple click of a mouse. AND its data recovery performance is *blistering* fast.

What more could you want from a backup product?

Space? O.K. Quick Restore takes up an astonishingly small amount of hard-disk space for index cataloging.

Quick Restore is simply the industry's most advanced, easy-to-use data backup & recovery solution for open systems workstations and file servers. Call 800.487.0080 today for more information and a free brochure.

Do the words *Easy-to-use, Speed, Innovative & Reliable* mean anything to you?

- ISO/POSIX Compliant
- Incredibly fast on-line indexing and volume location
- Open architecture supports multiple servers and DBMSs
- Data backup & recovery on a heterogeneous network without NFS
- Support for over 80 tape drives and robotics: 4mm, 8mm, DLT tape formats
- Exclusive compact database technology requires substantially less on-line disk space for catalog indexing



five overlook drive
amherst, nh 03031-2833
tel: (603) 672-8600
fax: (603) 672-3154

www.worksta.com

"Workstation Solutions' exceptional customer support organization makes it worth buying Quick Restore. Their staff is highly knowledgeable and efficient – it's a pleasure to call them!"



Don't let time pass you by! Get Quick Restore today.

Bazaar”) and me (long-term freeware groupie :-).

Tim invited us to meet so that we could explore the opportunities and problems that face the free software community. It was a heady experience for me and, I suspect, for many of the other attendees. Despite some digressions into favorite themes, we managed to discuss several substantial topics. What motivates developers to release their source code as free software? What characteristics set free software apart from proprietary code? What issues must the free software community solve if it's to be accepted by "Big Business"?

I won't try to detail all of the points that were made, but here are some that struck me as particularly interesting and meaningful:

- The terms "freeware," "free software" and "freely redistributable software" should be deprecated, in favor of a (trademarked) "open source" designation. (For what it's worth, I feel that open source is a special case of the other terms.)
- The characteristics of open source software should be explicitly defined. (A

draft definition is available at <http://www.opensource.org>) In short, however, open source software is open to inspection, use, modification and redistribution, within certain limits (such as the GPL restrictions), by any interested party.

- Massive independent peer review allows open source software to have far fewer long-standing bugs than proprietary software. Bugs can be found, analyzed and corrected by the most interested party (the user).
- The distributed development model used by open source software enables load sharing and diverse (evolutionary) development. It also eliminates a great deal of paperwork, reducing the transaction cost of cooperation.
- The open source model is useful for implementing certain forms of social action. The rapid spread of trustworthy cryptography, for instance, is centered in the open source community. (Don't trust the cryptography algorithm? Examine it!)
- Because the open source community has no proprietary bias, it tends to adopt and implement industry standards (for example, IPv6) with great speed.

Does the Boss Know?

Historically, open source software has been introduced in guerrilla fashion: Techies need to get a job done, know that a given package will help and put it in place. Only later does the boss find out that the package is in place. If the boss would learn a little bit more about open source software, however, he could gain far more.


Open source software is robust, flexible and safe from vendor policy changes. (Tried getting a copy of SunOS 4.1.4 recently?) It also promotes a climate of vendor-neutral standards, reducing the risk that any one vendor can "own" the entire playing field. As a businessman and an open source advocate, I find these arguments to be pretty compelling. ➔

Richard Morin operates Prime Time Freeware (info@ptf.com), which publishes mixed-media (book/CD-ROM) freeware collections. He also consults and writes on UNIX-related topics. He may be reached at Canta Forda Computer Laboratory, P.O. Box 1488, Pacifica, CA 94044 or by email at rdm@cfcl.com.


INDUSTRIAL STRENGTH Remote Access and Multiport Connectivity

Aurora multiport serial communications products provide high-performance remote access and WAN connectivity without the complexity or expense of Ethernet-based routers and servers. They offer the ideal solution for connecting modems, ISDN terminal adapters, DSU/CSUs, terminals, and instrumentation directly to a host system. Over 50,000 users worldwide, including leading telecommunications providers, financial services companies, and original equipment manufacturers, rely on Aurora for their serial connectivity requirements. Shouldn't you?

Contact Aurora for your free 42-page Guide to Multiport Serial Connectivity for Sun Solaris and Windows NT



- Available for Sun Solaris™ (SBus and PCI) and Windows NT® (PCI)
- Nine product families, 2 to 128 ports from one system slot
- Asynchronous connectivity up to 230.4 kbps
- WAN synchronous connectivity up to 256 kbps
- Support for RS-232, RS-422, RS-485, RS-449, EIA-530, and V.35
- Sun Solstice, Windows NT RAS, and SNA Server compatible



Phone: 781.290.4800 FAX: 781.290.4844
E-mail: info@auroratech.com
<http://www.auroratech.com>

AURORA
TECHNOLOGIES



Proper Searching

The other week I was searching the Web for help with a particularly perplexing problem on one of my Windows NT clients. "If only all the world were running Solaris!" I groaned while surfing my way to Microsoft Corp.'s technical support page. The Web has proven itself to be extremely useful in providing answers to difficult problems. A seasoned professional, I am more than happy to make use of other people's solutions and often search the Web before getting to the "RTFM" stage of problem-solving.

For my particular problem, little help was available. I drew a blank at <http://support.microsoft.com/support>. Every word in the free-form question that I posed was dismissed as noise by the search engine. Not used to making so little progress in my surfing, I began to think about the process of finding information, not only on the Internet, but in my own "back yard"—my file systems. How often have I forgotten the name of a file I was working on? Even with fairly

well-organized directories, I have lost things and wasted time looking for or re-creating them.

Genetically predisposed to be fond of regular expressions, I have sometimes run through my directories with a `grep` command designed to flush out files containing whatever keywords I remembered from my masterpiece in progress. Commands like

```
find /* -owner slee -exec \  
grep "netmasks" {} \;
```

filled my screen with sometimes helpful results.

Yet, while many of us are still "greeting" our way around our file systems, far more sophisticated search tools are not only available, but are part of our daily lives.

When I compose queries on Infoseek or HotBot, I'm not as constrained as I have been in most of my years as a UNIX sysadmin. Sure we have `egrep` and its siblings to help broaden our

search, but what tools in use on the Internet can we grab and use in our own file systems?

An Answer or Two

The first answer to my question came from memory. Years ago, I heard about `agrep`, a tool that is similar to `grep` but allows for "fuzzy" matching. The most significant feature that `agrep` provides that is not supported by the other members of the `grep` family is the ability to search for approximate patterns. For example, the command

```
agrep -2 netmasks myfile
```

will find "netmasks" as well as any other word that can be derived from it, given no more than two substitutions, insertions or deletions. In addition, `agrep` has a "best match" option, `-B`, that generates a list of matches with the closest (that is, those with the fewest substitutions) listed first.

Another difference between `agrep`

and the rest of the clan is that `agrep` is record-oriented rather than line-oriented. Although the default record *is* a line, `agrep` can be easily coaxed into considering portions of lines or groups of lines as a record. This is useful, for example, in pulling an email message from a user's inbox. The `-d` option is used to specify an alternate delimiter as shown here:

```
agrep -d '^From ' 'netmasks' mbox
```

This command displays messages containing the word "netmasks" from the mail box, `mbox`. The search pattern `^From` identifies the beginning of each message.

`agrep` also allows the user to specify multiple patterns along with Boolean (AND and OR) logic. The Boolean operator is expressed as a separator in the search string. A comma represents an OR. For example,

```
agrep -d '^From ' 'netmasks,subnet' mbox
```

displays messages with either of the keywords. A semicolon represents an AND. So the query,

```
agrep -d '^From ' 'netmasks;subnet' mbox
```

would only display messages with *both* keywords.

Clearly `agrep` is a useful addition to the `grep` family. It was developed by some very clever people at the University of Arizona Department of Computer Science (<http://www.cs.arizona.edu>) and appears to be in frequent use on the Web.

Glimpse – A Real Find

The second answer to my question came from surfing. In following links I came across while reading about `agrep`, I found numerous references to another popular tool—not surprisingly related to `agrep`—called Glimpse. This tool takes fuzzy matching from files to entire file systems.

A very powerful indexing and query system, Glimpse (or GLOBal IMPLICIT SEArch) allows you to search through all your files quickly and, as with `agrep`, use fuzzy logic. Because the tool uses an index prepared with the `glimpseindex` command (for example, `glimpseindex mydir`), you don't have to provide file names or directory names at the time of the search. Once the index is built, you can use much the same logic that you would use with `agrep`. The command

```
glimpse -2 netmasks
```

would find all of the occurrences of "netmasks"—with up to two substitutions, insertions or deletions. Running against the index, the command is fast. It doesn't have to read through the files in real time. Generally, the indexing operation will be run through `cron` overnight, incorporating new and changed files for the next day's searchers. Once you provide Glimpse with the word(s) that you want to match literally or fuzzily, it will provide a list of files containing the text or, if you prefer, display each line in the qualifying files containing the matches.

Glimpse produces very small indexes—as little as 2% the size of the indexed information—keeping its overhead to a minimum. If you prefer, you can elect to build a larger index for faster searching but, for most users, the smallest index seems fast enough. The largest index (called "medium") saves not only the files in which words appear, but their locations within the files.


Not completely foolproof, Glimpse doesn't work with a certain class of files—namely binaries and files containing irregular text like uuencoded data or MIME attachments. Files containing this type of data are ignored during the indexing and therefore aren't available for searching.

When setting up Glimpse on your system, keep in mind that the tool requires read permission for the index and all the files to be searched. This restriction may introduce complications if you have populations of users with different access privileges to the indexed files.

Glimpse is also a product of the University of Arizona Computer Science Department and is used by many search services on the Internet. If you frequently surf the Net, you may already be using it. If you'd like to see what it can do for your file systems, you can join the mailing list by typing "add" or "subscribe" in the body of an email message to `glimpse-request@cs.arizona.edu` (do the same with the word "remove" to get off the list).

Gee, I wonder what tool they're using to extract these keywords from their email. ➡

Internet and IT Costs Rising?



Minimize the Risk,
Rent or Lease a SPARC™ System.

Get Pre-Approved
Lease Line of Credit!

Rave Financial Services, Inc.



Fax: (810) 939-7431 • E-Mail: rent@rave.com
<http://www.rave.net>

MEMBER
ELA

1-800-500-7283

SPARC® is a registered trademark of SPARC International, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

S. Lee Henry lives about 50 feet off the shore in San Rafael, CA, and is always searching for something. If you'd like to tell her about one of your exciting discoveries, send email to slee@cpg.com.



LYNNE CANNOY

Sharing File Systems, Part 2

This month, we conclude our look at disks and file systems under Windows NT by considering a facility for sharing these resources with UNIX systems. The facility is called Samba, a name that plays on the acronym for the Server Message Block (SMB) protocol used to implement resource sharing in Microsoft Corp. networking. Samba provides UNIX systems with transparent access to Windows NT file systems, but capabilities for moving in the opposite direction are more limited. Network File System (NFS) is the usual UNIX method for making file systems available to other hosts across a network, and there are several good commercial products that provide NFS client and server capabilities for Windows NT.

Samba is freely available software distributed widely on the Internet. The official Samba Web site is in Australia (see <http://samba.anu.edu.au/pub/samba>). As of this writing, the current version, 1.9.18, is better than earlier ones. Its main advantage over the

immediately preceding version is that support for encrypted authentication has been folded into the main Samba release (previously, you had to acquire and build a separate encryption library and then rebuild Samba if you wanted this feature). In addition, some earlier versions of Samba (by now, several versions back) had significant security problems that are not an issue in the current version.

Samba features extensive documentation, including very good instructions for building it once you've downloaded the distribution. There is also a new book on Samba by John D. Blair, entitled *SAMBA: Integrating UNIX and Windows* (published by Specialized Systems Consultants Inc., Seattle, 1998, ISBN 1-57831-006-7).

What Is Samba?

Samba implements the SMB protocol for UNIX systems, allowing them to participate in native Windows networking and resource sharing. As we've noted, SMB is

the native protocol for Windows file and print sharing services, and Samba allows UNIX systems to participate in both arenas.

In addition, Samba can allow a UNIX system to function as a browse master, and the latest version includes alpha software support for providing some primary domain controller-related services.

SMB provides the services of part of the application layer and the presentation and session layers of the OSI Reference Model (see Figure 1). The

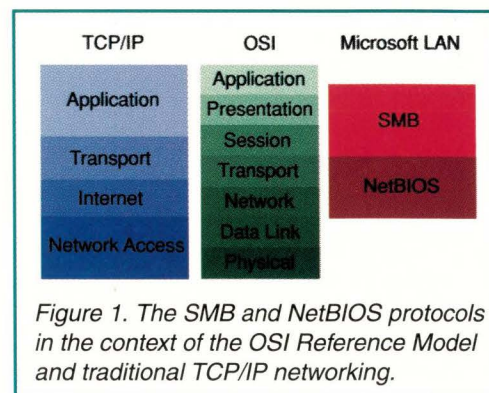


Figure 1. The SMB and NetBIOS protocols in the context of the OSI Reference Model and traditional TCP/IP networking.

NetBIOS protocol provides the functions of the transport and network layers.

As Figure 1 shows, these protocols require some lower-level networking capabilities to function. Traditionally, this was provided by NetBEUI, but in recent installations, SMB and NetBIOS usually run on top of TCP/IP.

Samba Configuration

Samba implements SMB via two daemons running on the UNIX system. The `smbd` daemon provides file system and printer sharing services, and the `nmbd` daemon handles NetBIOS name server requests. You have several options for running these daemons. The easiest is to run them as stand-alone servers, invoking them with their `-D` option, most likely from one of the system boot scripts.

Alternately, you can place them under the control of `inetd`, the TCP/IP master daemon, by modifying the `inetd.conf` and `services` configuration files (both located in `/etc`). In

this mode, they are started automatically by `inetd` on demand as requests for their services are received. The required configuration information is shown in Listing 1.

The Samba software is conventionally installed in `/usr/local/samba`. The configuration file that controls what local resources are made available through Samba (and to whom) is named `smb.conf` and is stored in the `lib` subdirectory of the Samba installation point (usually `/usr/local/samba/lib/smb.conf`). Listing 2 gives a simple Samba configuration file and includes annotations for the major entries within it.

The `[global]` section of the configuration file is used to define global settings that apply to all exports from the local UNIX system. In our example, we specify hosts that can and cannot access any of our exported resources, users or groups that are similarly allowed and denied access, and the maximum size of the Samba log file.

A section for an exported file system comes next in the sample configuration file, and its header line defines the share name

by which it will be known—in this case, `chemdir`, corresponding to the `/chem/data/new` subdirectory tree on the local system. Other entries in this section specify whether the file system is exported in read-only or read-write mode, users and groups that are allowed to read from it and write to it, and other parameters for the export. The `force group` entry deserves special mention. It logically maps all incoming user access to this file system to the UNIX group named `chemists`.

Listing 1. Sample inetd Configuration

Add to `/etc/inetd.conf`:

```
netbios-ssn stream tcp nowait root /usr/local/samba/bin/smbd smbd
netbios-ns dgram udp wait root /usr/local/samba/bin/nmbd nmbd
```

Add to `/etc/services`:

```
netbios-ssn 139/tcp
netbios-ns 137/udp
```

Listing 2. A Simple Samba Configuration File (Annotated)

```
[global]
hosts allow = vala, pele
hosts deny = lilith
valid users = dagmar, @chem, @phys, @bio, @geo
invalid users = root, admin, administrator
max log size = 2000
```

Global settings that apply to all exports.

```
[chemdir]
path = /chem/data/new
comment = New Data
read only = no
case sensitive = yes
force group = chemists
read list = dagmar, @chem, @phys
write list = @chem
```

*Define a directory (share) for export.
Local (UNIX) path to be shared.
Description of the file system.
File system is not read-only.
File names are case-sensitive.
Map all user access to this UNIX group.
Windows NT names allowed read access.
Windows NT names allowed write access.*

```
[homes]
comment = Home directories
writeable = yes
valid users = %S
```

Automatically define shares for each user name.

```
[laser4]
printable = yes
comment = LaserWriter on dalton
public = yes
postscript = yes
printer name = laz4
printer driver = driver name
```

Define a UNIX printer to share.

Specify the Windows NT standard driver name here

Listing 3. Use smbstatus to Monitor Shared Local Resources

```
Samba version 1.9.18
Service      uid      gid      pid      machine
-----
chemdir      nobody  chemists 14810    vala (193.0.13.4) Mon Jul 14 11:51:07

No locked files
```

Once a share has been defined in this way, it may be accessed by users on Windows NT systems by the UNC path or via the normal `net use` command. For example, the following command maps the `chemdir` share (located on the UNIX host named `dalton`) to drive S:

```
C:\> net use S: \\dalton\chemdir
```

The `[homes]` section of the configuration file automatically defines and sets defaults for shared user home directories, with the actual shares being created on the fly as each user requests access to it (a practice conceptually related to automounting on UNIX systems). Each user's share is given the same name as his user name. The `valid users` entry in this section illustrates the use of expandable parameters within the Samba configuration files. In this case, the share created for each user's home directory will be accessible only by that user as `%S` expands to the share name.

The final section of the configuration file makes a local printer available to network users using the standard Microsoft windowing printer sharing facilities. (A full discussion of sharing printers with Samba will be included in a future column.)

Samba Utilities

Samba includes several utilities that are useful for monitoring and maintaining a Samba installation. The `testparm` command may be used to verify syntax of the Samba configuration file. The `smbstatus` command displays information about current usage for shared local resources (see Listing 3). The `smbclient` command allows a UNIX user to access remote SMB shares, and it provides an FTP-like interface. The `smbrun`, `smbprint` and `smbtar` commands provide specific Samba client services for running a remote command, printing and executing a `tar` command, respectively.

Unfortunately, general transparent Windows NT file system access is not part of Samba at the moment. On Linux systems, the `smbmount` command allows a Windows NT file system to be mounted locally (supported by the `smbfs` facility in the kernel). On Solaris systems, the Solstice LM Server product provides similar functionality. On other UNIX systems, however, the client tools must be used to access remote shared resources.

Passwords and Samba

By default, Samba performs user authentication by sending user passwords over the network as clear text (like the FTP and Telnet facilities). Windows NT Service Pack 3 disables plain-text authentication and effectively "breaks" standard Samba.

You have two options for dealing with this issue. You can restore plain-text password support to Windows NT by setting the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Rdr\Parameters\EnablePlainTextPassword` to 1 (the data type is `REG_DWORD`).

However, a better alternative is to enable Samba support for encrypted passwords. The latter is accomplished by adding entries to the `[global]` section of the Samba configuration file, which vary according to what host you want to provide authentication services for Samba:

To authenticate users via the local Samba server:

```
encrypt passwords = yes
security = user
```

To authenticate users by a different Samba server:

```
encrypt passwords = yes
security = server
password server = host
```

Then, run the `mksmbpasswd.sh` shell script to create the initial Samba password file, which is what Samba uses to perform authentication. For example:

```
# cat /etc/passwd | \
/path/source/mksmbpasswd.sh \
> /usr/local/samba/private/smbpasswd
```


The created file must be owned by `root` and have mode 600. The protection mode on the private subdirectory must be 500.

At this point, the Samba passwords are identical to the UNIX passwords, but they are not automatically updated when the local password file changes. Instead, users must use the `smbpasswd` utility to change their Samba password. Earlier versions of Samba included a daemon that propagated Samba password changes back to the UNIX `passwd` file. Patch 4 to Version 1.9.18 provides an enhanced version of `smbpasswd`, which enables the command to update the Samba password, the UNIX password and the Windows NT password in a single operation.

In conclusion, Samba provides powerful and reliable file system and printer sharing services from UNIX systems to Windows NT and other systems across a Microsoft network. It is yet another testimonial to the excellent software coming out of the free software movement and to the skill and generosity of these programmers. ➡

Aleen Frisch is systems administrator for a very heterogeneous network of UNIX and NT systems. She is also the author of the books Essential System Administration and Essential Windows NT System Administration (both from O'Reilly & Associates Inc.). In her (almost nonexistent) spare time, she enjoys painting and lounging around with her cats, Daphne, Susan, Talia and Lyta. Email: aefrisch@lorentzian.com.

- ▲ wizard's apprentice
- ▲▲ super user
- ▲▲▲ wizard



Surmounting Unmount Woes

Q: I am sometimes unable to unmount a file system. Unmount continues to say the device is busy, even though I run `fuser` on the file system and it does not find any open files. Why can't I unmount the file system? Are there open files that `fuser` doesn't show? ▲▲▲

Bob Richter,
R&R Software Inc.

A: You're correct that you cannot unmount a file system if any of the files on that file system are open. Also, you cannot unmount if a user's current directory is on the file system. But even if you ask all your users to log out, you still may not be able to unmount a file system. This is probably due to a program running out of that file system. For example, if the program `/tmp/a.out` is running, you would not be able to unmount `/tmp`.

There are two common tools that help you find open files: `fuser` and `lsprof`. `fuser` will not find those exe-

cutables; in some circumstances `lsprof` will. Let's first check out these tools, then we'll look at another method for finding those wayward executables that are keeping you from unmounting your file system.

fuser: Easy to Use

`fuser` is the standard AIX program for finding open files. It comes with AIX and is easy to use. To list the IDs of all processes with a particular file open, type

```
#fuser -u filename
```

The `u` parameter appends the process's owner to each process ID.

If you want to find all processes with files open on a particular file system, you have to use the block device name for

that file system. Use `df` to find the block device name, as shown in Figure 1.

When you run `fuser` on the block device—`/dev/hd9`, in this example—you will get information on all open files in the entire file system. `fuser` will also find open files on NFS-mounted file systems.

`fuser` will not find running executables. If, for example, the program `/myfs/a.out` is executing, you cannot unmount the `myfs` file system. But, `fuser` will not find `/myfs/a.out`; you need a better tool.


lsprof: A General-Purpose Tool

`lsprof` is a general-purpose, portable open-file finder written by Victor Abell of the Purdue University Computing Center. You can get a copy of `lsprof` by anonymous FTP to `vic.cc.purdue`.

Figure 1. Sample df Command

```
# df /myfs
Filesystem      512-blocks    Free   %Used   Iused   %Iused   Mounted on
/dev/hd9          131072    86816    34%    1472     5%      /myfs
```


MIS DEPARTMENT



BANG HEAD HERE!



Aspen takes the
pain out of
System
Integration!

Aspen
CONSULTING INC

IBM® ADSM CERTIFIED • IBM® HACMP CERTIFIED • IBM® RS/6000™ SP™ CERTIFIED

888-249-4968 • info@asp.com • www.asp.com • FAX: 847-806-1260

An IBM Premier Business Partner



The IBM logo is a registered trademark and the IBM Business Partner Emblem is a trademark of the International Business Machines Corporation and are used together under license. AIX is a registered trademark and RS/6000 and SP are trademarks of International Business Machines Corporation. All trademarks are property of their respective companies.



edu. Both sources and precompiled binaries are found in /pub/tools/unix/lsof/. Due to the program's intimate dependency on the specifics of your operating system, it is recommended that you compile and build lsof on your own system. Get the latest tar file—Version 4.29 as of this writing—and untar it. In the source directory, run ./Configure, then make. If you want to let unprivileged users have access to lsof, you may have to change some installation options. To find who has a particular file open, just run lsof on the file, as shown in Figure 2. The general form is:

```
lsof [options] path
```

Note especially the following useful options:

- -help displays a summary of lsof usage.
- -f forces lsof to interpret the path name as a file. By default, if the path name is the mount point of a file system or is a block device, lsof assumes you want to see all open files on that file system.
- +f forces lsof to interpret the path name as a file system. If it's not a file system mount point, you get an error.
- -i i selects files matching a selected Internet address.

Figure 2. Run lsof to Find the File's Owner

```
# lsof /myfs/dat.txt
COMMAND  PID  USER  FD  TYPE  DEVICE  SIZE/OFF  INODE  NAME
tail     10562  root   0r   vreg   10,7     2469     42    /myfs/dat.txt
```

Figure 3. Run ls and df to Find Major/Minor Numbers

```
# df /myfs
Filesystem 512-blocks  Free  %Used  Iused  %Used  Mounted on
/dev/hd9    131072    86816  34%    1472    5%    /myfs

# ls -l /dev/hd9
brw-rw---- 1 root  system  20, 8 Aug 26 1997 /dev/hd9
```

- -p p selects open files for the specified process.
- -u u selects open files for the specified user. User can be either user IDs or user names.
- -X finds shared libraries and executing programs. This is the option that will find those open-file problems fuser skips over. However, there's a caveat: On busy systems, the -X option can cause the system to hang. This is a known but unfixed AIX bug. I've never encountered the problem myself, but you should use -X only when you have to.

crash Can Find Open Executables

If you think an executing program might be keeping your file system unmountable, and you can't find it with lsof, try crash. No, that's not what will happen to your system.


ALPHANUMERIC PAGING FOR UNIX

RELIABLE, EASY DELIVERY OF MESSAGES ANYTIME ANYWHERE

- Email forwarded to pager automatically
- Pages can be generated from scripts, and network monitoring programs
- GUI and command line interface
- Works with any paging service
- Automatic email confirmation, history logs and error reporting
- Client-server technology
- Works with digital and alphanumeric pagers

Personal Productivity Tools for the Unix Desktop

14141 Miranda Rd
Los Altos Hills, CA 94022
Email: sales@ppt.com
Tel: (650) 917-7000
Fax: (650) 917-7010
<http://www.ppt.com>



Countless calls to faceless people = wasted valuable time.

Especially when you can make one call for a total solution to your IT needs. At One Source, we know that you may have issues such as how to seamlessly integrate an existing PC LAN with robust UNIX enterprise servers. Some clients need to concentrate on finding a mid-range solution while others have a basic concern about the best hardware to buy — IBM, HP, SUN, DEC. You may find that you have software integration issues—Java, Oracle, middleware, etc. And what about internet/intranet deployment? Don't panic. Make one call to One Source.

One Source Computer Corporation Your one source IT solutions provider

- UNIX Training, Consulting and Systems Integration Services
- Sun Master Services Agreement (MSA) Supplier and Catalyst Developer Program Partner
- Oracle Alliance Partner
- Full-Service Authorized Reseller of IBM and DEC Hardware and Peripherals
- Sun and HP Refurbished Equipment



One Source - the single source to supply you with trained consultants to analyze your needs and provide solutions for installation, custom development, training, and mentoring when you need it.

119 Herbert Street Framingham, MA 01702
<http://www.onesourcecorp.com>
e-mail: onesourcecorp.com 508-370-3960

crash is the AIX system tool for examining core dumps, but it will also work on a running system. Here's how to use crash to find the open executable. First use `ls` and `df` to find the major and minor numbers of your file system (see Figure 3). The major and minor numbers in this example are 20, 8.

Next, run `crash`. It's interactive. You'll first see:

```
# crash
WARNING: Using crash on a live system can
potentially cause a system crash
and/or data corruption.
```

That warning message should encourage you to be careful. But you can run `crash` on a live system without crashing it.

Use the `inode 0` command to show the column headers: ADDRESS, MAJ, MIN, INUMB, REF, LINK, UID, GID, SIZE, MODE, SMAJ, SMIN and FLAGS. The `inode` command without parameters will display a lot of inodes. You may want to pipe that list through `awk` to select just your file system. Do

Figure 4. Pipe List through AIX awk to Select Your File System

```
> inode | awk '{if ($2==10 && $3==8 && $5>0) print $0}'
0x09e2b340 010 0008 16 1 2 100 100 512 d---755
0x09e2b490 010 0008 20 1 1 0 0 3521 f---755
```

Figure 5. Use find to Find the File's Name

```
# find /myfs -ls | grep -w 20
20 4 -rwxr-xr-x 1 root system 3521 Mar 30 15:43 /myfs/a.out
# find /myfs -ls | grep -w 16
16 1 drwxr-xr-x 2 100 usr 512 Mar 30 15:43 /myfs
```

that by selecting the major and minor numbers and a REF count greater than 0 (meaning the file is in use), see the example in Figure 4.

The fourth column is the inode number of the open file. Use the `find` command to find the file's name, as shown in Figure 5. There's the program that is keeping the file system busy.

Do keep `crash` in mind, but usually `lsOf` is the only tool you'll need. Don't leave home without it. ➔

Jim Fox works as a systems programmer for the University of Washington. He writes and maintains distributed applications that run on a variety of UNIX systems—and some non-UNIX ones. He is also the deputy manager for the Interoperability Project for SHARE's Open Systems Group. Email: fox@cac.washington.edu.



Blow Away Database Speed Limits!

The fast track to database performance begins and ends with MegaRam® Solid-State Disks (SSD). These plug and play SCSI devices shatter I/O bottlenecks and do more to improve database response than faster processors and more powerful database engines combined.

It takes only a small amount of MegaRam SSD to accelerate your database response by 100% to 400% because most I/O traffic is caused by about only 3% of your files. In Oracle OLTP environments, the *redo log* and *rollback segments* are the "hottest" files. By placing these and other hot files—*system tablespaces, indices, and temp files*—on a MegaRam SSD we guarantee significant performance improvements!

Let Imperial analyze your system for I/O bottlenecks. MegaRam Solid-State Disks save you time, money and aggravation, and above all they remove the roadblocks to blistering database speed.



MegaRam SSD's are platform and software independent, scalable, ultra-reliable, and affordable. Capacities range from 32 megabytes to 8 gigabytes in a single chassis.

Start accelerating your data today:

800-451-0666

Visit our Website at:

<http://www.imperialtech.com>

IMPERIAL

Imperial Technology, Inc.

2305 Utah Avenue • El Segundo, CA 90245
Telephone: (310) 536-0018 • FAX: (310) 536-0124



Business Alliance Programmer



SOLARIS COMPATIBLE



HEWLETT PACKARD

Circle No. 22

SunExpert Magazine ■ June 1998

Datagrams

by John S. Quarterman



Processing Masses of Mail

A few months ago, I wrote a column entitled “Essential Tools” (October 1997, Page 78) and later “More Network Tools” (January 1998, Page 62). One of these mostly unsung tools is `procmail`, which permits you to filter incoming mail and file messages to a file, remail them or let them be delivered normally. It’s easily worth a whole column to itself, so here it is.

Too Much Mail

Says one systems administrator: “Even in a bad month like last month [March 1998], when I wasn’t filtering as much stuff as I might like to, I threw away 3,314 messages to `/dev/null`, with an average size of 86,853 bytes. I accepted in my mailbox 8,505 pieces of mail with an average size of 4,844 bytes. Of those 8,505 messages, there were 866 from an external support list, 2,873 from my own group’s support list and 138 gripe list messages. The rest were mostly

individual pieces of nonlist mail.”

She continues: “I file many messages into mailboxes for later perusal. That month, I filed 105 from the SAGE members’ list, 20 from the SAGE jobs offered list and 30 from a dorm support list. I also filed 5,837 bounce messages, which I normally would have sent to `/dev/null/` since other people usually handle postmaster mail, but I was looking for something specific. And 16 requests to subscribe to a mailing list that I’m the owner of. I usually look at those files at the end of the day, so I can handle more urgent messages first.”

The previous month (February 1998), she sent 5,443 messages to `/dev/null`, with an average size of 455,906 bytes. She accepted 8,134 messages that month.

In January 1998, she accepted only 7,514 messages. This was because she was sick and not answering her mail rapidly—when you don’t answer your mail, you don’t get as much mail.

Curiously enough, that averages out to about 280 messages per day, which is very close to the 300 messages per day that the Message Handler (MH) mail user interface documentation brags that MH can handle. If it takes on average 30 seconds to handle each message, that amounts to 150 minutes, or 2.5 hours, which is a lot of one’s day.

In other words, even with efficient mail-handling software, there is still a human processing limit on the amount of mail that a single person can handle in one day.

Says our example systems administrator: “I get way too much mail.” I’m sure many of the rest of us feel the same way. It would be much worse without `procmail`. Instead of seeing 8,505 messages directly in her mailbox, she would have seen more like 18,000, or 600 per day. That would be overload.

To use `procmail` effectively, you need to tell it what to filter. You do this by creating a `.procmailrc` file in your home directory.

Datagrams

Here is an example `.procmailrc` file:

```
PATH=/usr/local/bin:/bin:/usr/bin:
MAILDIR=/home/username/Mail
#DEFAULT=$MAILDIR/mbox
LOGFILE=$MAILDIR/from

:0
* ^Subject:.*Output from.*
${MAILDIR}/in/cron

MYTO = "(To|Cc|Resent-?(To|Cc)):"

:0
* ${MYTO}.mudpp-list@spice.com
/dev/null

:0
* ^Reply-To:.*list@studiob.com
! r-cbpl@zilker.net
```

This file first sets up some parameters. The search path, `PATH`, is set up with its standard UNIX syntax. `MAILDIR` is a variable that is used explicitly in later rules. It contains the path name of the directory where you save your mail. (Replace `username` with your user name.) We'll discuss `LOGFILE` later.

Each `:0` is the beginning of a rule, which ends with the next blank line. In between are patterns and an action. The first rule in the example says all mail with `Subject:` containing the string `Output from` should be filed in `${MAILDIR}/in/cron`. The wild cards use regular expression syntax like `sed` (or `awk` or `perl`), not like filename wild cards. To match multiple arbitrary characters, you must say `.*`, not just `*`. Then we define a variable, `MYTO`, to match several header tags, including `To:` and `Cc:`. We use `MYTO` in a rule that says to look for all mail with any header with one of those tags and a certain address in that header line and to file that mail in `/dev/null`, thus throwing it away sight unseen.

The final command in the example sends mail from a certain mailing list to another host. The exclamation point at the beginning of the line tells `procmail` to remail the message to the address on that line, rather than trying to save the message in a file. By default, when `procmail` sees mail that matches no rule to tell it to do something else, `procmail` will deliver it to the usual place, such as `/var/mail/username` or `/var/spool/mail/username`. `procmail` is a mail delivery agent and can be used as such by `sendmail` or whatever other mail transfer agent that is in use on your system. You still read your mail using whatever mail user agent you usually use, such as `Pine`, `Elm`, `Mail` or `MH`.

When you have a `.procmailrc` file drafted, you can turn on `procmail` by editing your `.forward` file to contain a line like this:

```
"|/usr/local/bin/procmail"
```

The exact path name depends on where `procmail` is installed

on your system. As for installing `procmail` itself, I refer you to its documentation, which is complete and straightforward.

How can you achieve statistics like those quoted above? First tell `procmail` to gather data, by putting a line like the following in your `.procmailrc` file:

```
LOGFILE=$MAILDIR/from
```

Then, periodically run `mailstat`, which comes with `procmail`, on that file. Many people run `mailstat` out of their `crontab` once a month. Beware: `mailstat` re-names the data file, in this case, `from` to `from.old`. So if you run `mailstat` twice, you will lose the data. Save the file under another name first if you want to keep the data.

An Effective Antispam Device


Here's a way to get rid of some typical spam:

```
TRASHSUB="MAKE MONEY FAST!!\
|FREE CASH GRANTS\
|A FAST AND EASY WAY TO MAKE MONEY"

* ^Subject: ($TRASHSUB)
/dev/null
```

This `.procmailrc` fragment defines a variable with several subjects, and then uses a single rule to throw away all mail with

Why pay more?



{For the same sparkle.}

Our memory is designed to Sun specifications, built with the same components and tested for the same performance. We build 100% compatibility into every module. We're so dedicated to quality, we offer a lifetime warranty.

To get the same sparkle for less, ask for Transcend. The only people who see the difference will be from finance.

Transcend
The Memory Specialist

800-886-5590
www.transcendusa.com

Phone: 714-921-2000
Fax: 714-921-2111

Transcend Information, Inc.
1645 N. Brian St., Orange, CA 92867



Datagrams

those subjects. You can add more subjects easily, thus filtering still more spam. You can also define similar lines to reject any mail with `friend@public.com` in the `From:` header.

`procmail` can even filter on the body of the message. This would permit automatically bouncing "E-MAIL MARKETING WORKS!" and similar dreck. Personally, I prefer to send a complaint about each piece of spam I receive back to the originating or transit system's abuse address (see <http://www.mids.org/nospam/>, for more on this). I do use a script to automate composing such complaint messages, but I send them manually; I don't have `procmail` do it because I don't want spam automatically causing messages to go out from my machine. That would risk spam causing spam.

Ever get mail from somebody who just won't quit? Who continues to tell you about his latest solution for saving the world through shark bile? Or who is convinced some offhand comment you made is a dire insult to her affinity group? File it to `/dev/null`! No more bozo.

Other Uses

`procmail` can also be used to throw away any message that's bigger than about 5 MB. This eliminates almost any message with attachments in Microsoft Corp. Word format. Anyone who wants to communicate with a UNIX systems administrator can learn to use a portable format. I can't impose that limit, because I do sometimes need to receive GIF or JPEG images as attachments for use in Web pages and the like.

Another popular use of `procmail` is as a vacation message responder. That's a little more complicated than I want to discuss in this column. However, like many uses of `procmail`, you can borrow a rule from somebody else and not have to figure it out from scratch yourself. Now if I can just figure out how to get it to make coffee, it will be perfect.

How to Get It

If `procmail` is not installed on your system, you will need to retrieve its source code. Look first for a USENET archive of the newsgroup `comp.sources.misc`. If you can't find that, try `ftp://ftp.informatik.rwth-aachen.de/pub/packages/procmail/procmail.tar.gz`. A tip of the Datagrams hat to its author, Stephen R. van den Berg at RWTH-Aachen, Germany, `berg@pool.informatik.rwth-aachen.de`.

In conclusion, `procmail` is versatile, powerful and indispensable. →

John S. Quarterman is president of Matrix Information & Directory Services Inc. (MIDS), which publishes Matrix Maps Quarterly, Matrix News (monthly) and the MIDS Internet Weather Report (daily). John has written or co-authored seven books, but the best known one is still The Matrix. For more information, see <http://www.mids.org>. He can be reached by email at `jsq@mids.org`, by voice at (512) 451-7602 or by fax at (512) 452-0127.

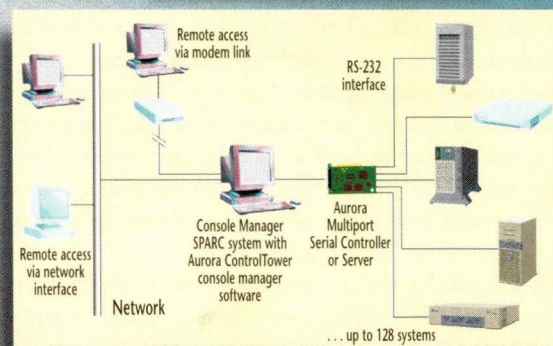

Aurora ControlTower™ CONSOLE MANAGER SOFTWARE

puts you in control of your Sun servers — from anywhere

Aurora ControlTower console manager software takes you beyond the limitations of hardware switches and network-based tools. One SPARC Solaris server functions as a common monitor and keyboard for up to 128 SPARC systems. But that's only the beginning. ControlTower software utilizes a serial console connection so you can access managed systems, perform diagnostics, and reboot a system even when the network is down. Remote access capabilities let you respond to a system emergency anywhere in the world, from office or home.

- Monitor and control one or more data centers from a central or remote location
- Same functionality as a keyboard attached directly to a system
- Access real-time data and a time-stamped log of system messages
- User-definable parameters for all or any managed system
- Security features control access to console manager servers and/or managed systems

For information on how to obtain a free evaluation copy of ControlTower console manager software, go to <http://www.auroratech.com>.



AURORA
TECHNOLOGIES
A Carlo Gavazzi Group Company

Circle No. 24

Phone: 781.290.4800 FAX: 781.290.4844
E-mail: info@auroratech.com
<http://www.auroratech.com>

What makes SharkRacks the most advanced racking solution on the planet?



SharkRack A5000


SharkRack has developed the most advanced racking technology in the world. SharkRack represents a paradigm shift in racking development with unprecedented R&D testing and co-development at Sun Microsystems's labs in the Silicon Valley. SharkRack is the **ONLY** racking manufacturer to be tested at Sun's labs, as well as the **ONLY** racking manufacturer to pass Sun's rigorous shock and vibration testing (on our first try!).



SharkRack E3000

The testing success realized by SharkRack was due to two extremely important factors; Thermal and Structural.

**For quotes and more info go to:
www.sharkrack.com**




SharkRack E4000



SharkCAB E4000 Cluster



SharkCAB E3000 Cluster




SharkKEY

Circle No. 36



CATALYST DEVELOPER



harkSTOR 112/114



DANIEL O'CONNOR

A Picture Worth a Million Bits

A few years ago, I purchased my first digital image scanner. It was a small handheld model and one of the first to hit the market for home and small business use. At the time, I thought it was the coolest device one could add to their private horde of workstation paraphernalia. I started scanning and showing off every snapshot in my photo collection—much to the chagrin of family and friends.

Everything was fine until I tried my hand at an 8-by-10 inch picture that was wider than my scanner. Fitting together multiple scan passes on large photographs is an acquired skill—or in my case, a quickly acquired headache.

Technology continued to advance and prices began to drop at a rate slower than the erosion of my patience. But I was soon able to replace the handheld scanner with a new flatbed scanner. Life was good again and I resumed scanning photos for my Web site. Always a mal-content, I soon began to yearn for a faster and more consistent method of updating

the image content on my Web pages. There had to be a way to streamline the time-consuming cycle of composing photographs, waiting for the film to develop and scanning the photos into digital format.

Enter the electronic digital camera.

Digital vs. Film

Digital camera technology has come a long way in the past few years. Improvements in image resolution, encoding, storage, optical features and convenience have made them an attractive alternative to traditional film-based cameras.

Although digital camera technology is not quite ready to completely replace the film-based camera, it does offer many features that make it a better fit in situations where electronic format is a premium.

But before you decide to throw away your 35mm and rush out and buy a digital camera, make sure you understand all the pros and cons of digital versus film formats. You might find that you'll need

one of each in your photographic arsenal.

Film provides the best overall image quality at a lower cost compared with digital media. This is because the analog nature of traditional film provides for a high level of control over the detail of the captured image. Each frame of silver-halide film contains millions of overlapping crystals that react chemically to variations in light intensity.

In contrast, consumer-grade digital cameras capture images using a matrix of a few hundred thousand to a million discrete pixels. Pictures created using film can be easily enlarged without incurring the distortion caused by pixelation inherent in a digital format. Information lost in the gaps between the pixel elements of digital devices becomes visible as jagged edges and "color noise" when the image is enlarged. More about this later.

There are professional-quality digital cameras available that incorporate much higher resolution-pixel matrices but not without an almost exponential increase in cost. If you want large, high-quality hard

copy at low cost, then stick with film. Choose a digital camera when you need the convenience of an electronic format. No more waiting for the film to be developed only to find that you took one good shot out of 24. Most digital cameras allow you to view the images while they are still in the camera. If you don't like a particular picture composition, then delete it and try again. Once you're satisfied, you can upload the pictures onto a computer, add them to Web pages, send them as email attachments to Grandma or print them using a photo-quality printer.

You can also use digital image-processing software to manipulate the composition or correct anomalies in the image. These features make digital cameras ideal for journalists, real estate brokers, insurance adjusters, Web developers, electronic artists or even amateur photographers. They're no longer just for those coffee mug and T-shirt shops in the local shopping mall.

Once you've decided to buy a digital camera, you need to carefully consider what features are most important to your particular image creation process and output medium. For example, if you're more the "point-and-shoot" type, you might prefer auto or fixed focus over manual focus. If the final product is going to be a Web page image, you may be able to avoid the higher cost of high-resolution cameras.

Picture Resolution

Most digital cameras currently on the market use a Charged Coupled Device (CCD) chip as the digital counterpart to film in traditional cameras. A competing technology to CCD is the Complementary Metal Oxide Semiconductor (CMOS) chip. CMOS chips are cheaper to bulk produce than CCD chips, but they are not as sensitive to light due to a smaller pixel area and they produce more color noise than CCD. Improvements in CMOS chip design could address some of these limitations in the near future.

A simple silicon CCD element comprises a matrix of light-sensitive pixels that react electrically to light intensity. Variations in the level of incident light striking a pixel are represented as a discrete gray scale from 0 (darkest)

to 255 (brightest). To incorporate color information, groups of three pixels are combined with red, green and blue filters to represent 16 million color variations. Under this simple architecture, only one-third of the total number of pixels are used to represent an image because of color filtering. There are more advanced color encoding implementations, but it's important

to keep this in mind as we look at the interaction between CCD resolution and picture size.

Low-end digital cameras use CCD chips with a matrix resolution of 350,000 pixels. Intermediate devices run in the 850,000- to 1.2 million-pixel range. High-end professional cameras have matrix resolutions exceeding 6 million pixels.


Serial Port Simplicity

For products that are simple to install, UNIX & NT drivers that are easy to load, and ports that administer just like on-board ports, look no further than Central Data. But don't let the simplicity fool you. Our solutions are truly innovative, and super fast.

EtherLite® Port Servers – Real tty or COM ports, located anywhere on your network. Compatible with 100MB Ethernet.

scsiTerminal Servers™ – Up to 32 ports from a single SCSI ID. No card slots needed.

PCI Serial Cards – 4- and 8-port models available.



To simplify your serial communications, call Central Data today at 1-800-482-0398.

Central Data

1-800-482-0398
www.cd.com
info@cd.com
217-359-8010
fax 217-359-6904

Now available
through SunExpress



1-800-USE-SUNX

Copyright © 1997 Central Data Corporation. All rights reserved. EtherLite is a registered trademark and scsiTerminal Server is a trademark of Central Data Corp. UNIX is a registered trademark in the U.S. and other countries, licensed exclusively through X/Open Company, Ltd. Windows NT is a registered trademark of Microsoft Corporation.

Circle No. 25

CCD technology is also used in other imaging systems such as telescopes, scanners and bar-code readers.

Picture Size

Most digital cameras provide selectable resolution levels that allow you to make trade-offs between image storage and image quality. At the high-end, the image resolution is usually close to the pixel-matrix resolution of the CCD or CMOS chip. For example, a camera might support an upper limit of a 1,024-by-768-pixel image from a 850,000-pixel chip, yet also offer lower resolutions like 640 by 480 (see Table 1). Remember that you end up sacrificing some CCD pixels to support color.

Some cameras offer maximum resolutions that exceed the pixel-matrix size of the CCD or CMOS chip. Software interpolation and compression algorithms are used to approximate missing pixel information. These techniques can introduce noise or artifacts into the image. Dark areas within an image are subject to higher levels of artifacts because the camera has to guess whether they are actual dark areas or whether particular wavelengths of light have been missed due to color filtering. The result is akin to the graininess in images produced with high-speed film (with large silver-halide crystals). Information is also lost if the image is compressed to save storage space. Artifacts become problematic as you try to enlarge the image for display on a computer monitor or for printed hard copy.

For "photorealistic" picture sharpness and definition, you need to maintain a minimum resolution of 300 pixels per inch (ppi). This usually isn't a problem when displaying images on a computer monitor because most common display devices range somewhere between 70 to 100 ppi. A 320-by-240-pixel image will look fine on a computer screen, and it will require less storage space than if it were to be displayed at a resolution of 1,024-by-768 pixels, assuming the number of colors is consistent.

Printers are another story. Most common hard copy devices provide resolutions of 300 ppi, 600 ppi or higher. At 300 ppi, a 3-by-5-inch print is going to require a 900-by-1,500 base image resolution to minimize visible artifacting.

There are a number of photo printers available on the market that incorporate additional software interpolation algorithms and hardware subsystems to enhance the sharpness of images printed from digital cameras.

Bottom line: If you want to print high-quality photos, you are going to need a camera with a CCD resolution of 850,000 or better. If your output medium is a Web page, then a low-resolution, low-cost camera will do fine.

Table 1. Image-to-CCD Resolution

Resolution	CCD
640 by 480	350K
1,024 by 760	850K
1,280 by 960	1.2M

Table 2. Common Media Storage Capacities

Media Type	Storage Capacity
Flash memory cards	256-160 MB
Smart media/Solid State Floppy Disk Cards (SSFDC)	2-8 MB
Compact flash cards	4-32 MB
Hard drive cards	128 MB-1 GB
Floppy disks	1.4-2.8 MB

Focus and Flash

Just like its film-based counterpart, lens configuration and exposure control are crucial to the quality of the images taken by a digital camera. Let's start with the space used to compose a picture, the viewfinder.

Most digital cameras come with a small LCD panel for viewing the image composition. Cameras with hinged LCD panels are great for taking pictures at difficult angles that would otherwise strain your neck. The addition of an optical viewfinder is helpful in situations where lighting can make the LCD difficult to read. Whether optical or LCD, make sure you know whether or not there is any differentiation between the image you see in the viewfinder and the image the camera lens sees; they are not always the same. A good single-lens reflex (SLR) configuration with a through-the-lens

(TTL) viewfinder will take care of any discrepancies between how the lens and viewfinder frame a scene.

Next, consider the lens configuration. Is the lens plastic or glass? Generally, glass will provide a clear image and do a better job of gathering light.

Is zoom provided, and if so, is it continuous or digital? Digital zoom uses discrete steps to increase magnification rather than a smooth continuous expansion of the scene. Common digital zoom systems will step from 1x, 1.5x, 2x, 2.5x to 3x. Is the focus fixed, auto or manual? Your choice will depend on how you intend to use the camera. A nice feature on some cameras is a threaded lens ring for attaching external filters, shades and additional lenses.

Finally, how much control does the camera provide over exposure? Many cameras allow you to override the auto-exposure settings by increasing or decreasing the "exposure correction" control. If the camera includes a flash, make sure you can turn the flash off in low-light settings or where flashes are prohibited—such as inside museums or churches. Digital cameras perform quite well without a flash in low light. Compare the International Standards Organization (ISO) or American Standards Association (ASA) ratings of a camera. A camera with a higher ISO or ASA number will perform better in low light. Because you can immediately view the image, you have the opportunity to play with exposure settings to fine-tune the final picture. In situations where you need a flash, pay attention to how long it takes the flash charge to recycle. This is a concern whether the camera is digital or film based.

Another nice feature of digital cameras is the ability to reuse internal and removable storage. If you take a lot of pictures, then you are spending a good deal of money on film and developing. Removable storage media simplifies the task of moving images between the camera and the computer. In many cases, you may be able to use the same storage media for your digital camera, palmtop, laptop and your desktop computer. Some common storage media types and capacities are listed in Table 2. Most of these devices are Advanced Technology Attachment

(ATA)-compliant and can be accessed from a computer using a PCMCIA adapter or external player.

Image Transfer and Manipulation

Most cameras provide serial and/or parallel interfaces for uploading images to the computer. Twain drivers are provided to interface with most popular image-manipulation software. Some of the cameras also incorporate National Television Standards Committee (NTSC) and Phase Alternating Line (PAL) interfaces, which can be used to plug the camera directly into a standard TV or VCR.

Another nice feature of digital image format is composition can be manipulated using software control on a computer. Tools such as Adobe Systems Inc.'s PhotoShop or GNU's GIMP can be used to correct artifacts and distortions, whiten flash-red eyes or completely change the content of the image by altering the background and applying special filters. This is where the real fun begins. Why not put a funny hat and a big nose on that picture of your boss.

You may find it helpful to select a camera that stores images in the same format you intend to use in the final product. This way you limit the loss of image information in the conversion from one format into another. For example, if you intend to use JPEG images on your Web site, then look for a camera that supports JPEG. Note that some cameras use proprietary image formats.

There's far more to digital photography than I've covered here. I don't pretend to be a professional photographer, just a Web site administrator who is tired of traditional film developing and scanning. Seek advice from a real expert if you're interested in digital photography. I've listed a few places to start your search in Table 3. You will also find useful information at many of the online computer and camera super-


stores. Do your homework and you'll find a camera that works best for you.

After all, you're going to want to make that Pinocchio picture of your boss look real good. Remember to go high resolution for those large copies to post on the office bulletin board. →

Jim DeRoest has been involved (for better or worse) with IBM UNIX offerings from

the IX/370 days, through PC/IX, AIX RT, AIX PS/2, AIX/370, PAIX, AIX/ESA and AIX V3. He is employed as an assistant director supporting academic and research computing at the University of Washington, and is the author of AIX for RS/6000—System and Administration Guide (McGraw-Hill). He plays a mean set of drums for the country gospel band Return. Email: deroest@cac.washington.edu.

RAVE COMPUTER ASSOCIATION, INC.



Doing the extraordinary isn't easy, but we've become pretty adept at it. We have been meeting and exceeding our customers critical computing needs for over ten years.

Armed with a terrific attitude, our sales team continues to win the respect of our customers. We stand behind our products and services.

Let Rave be the choice for your future technology needs.

**Rave Desktop 5 ❖ Rave Desktop 20 ❖ Rave Desktop U2
Rave RackMount U2 ❖ Axi Technology
Rave Power Tower ❖ Rave Power Rack**




**RAVE COMPUTER ASSOCIATION, INC.
Call Today at 1-800-966-7283**

**36960 Metro Court, Sterling Heights, MI 48312
Visit our web site: www.rave.net**

Table 3. Useful Digital Camera Information Sites

- <http://www.steves-digicams.com>
- <http://www.pcphotoforum.com>
- <http://the-tech.mit.edu/KPT/KPT.html>
- <news://rec.photo.digital> (newsgroup)



*Sir, I have found you
an argument; but I am
not obliged to find you
an understanding.*
– Samuel Johnson

We Argue with Our Code

There are several ways to get information into a UNIX program to affect the way it behaves. In this column we're going to talk about one of these: command-line arguments.

First, though, let's briefly list all the ways. There are four:

1. **The environment** – Each UNIX process inherits a suite of information from its parent process, which it can interrogate. All the programs you run from the shell have a collection of state information, including the current working directory and the current user's ID, which the operating system keeps track of during the life of the process.

These values are private to the process. Once the process has started running, it can manipulate most of them, but you can't change them from outside. The fact that you can't change an environment from the outside is very important. In all the programming newsgroups, you'll find a never-ending stream of posts asking, "How do I change the value of an environment variable in a parent process

from one of its children?" The answer (always found in the FAQ) is, "You can't."

Nevertheless, if you tweak the environment before you start a program, it can examine that environment at start-up time and decide how you want it to run. Most of this state information is held in environment variables, which you can see with the command `env`.

We could say a lot more about the environment, but we won't.

2. **Interprocess communication** – Running processes can talk directly to one another in a host of interesting ways, collectively called "IPCs." The most common of these are pipes (named and unnamed), signals, sockets, semaphores, message-passing and shared memory. Not all vendors supply all of these, but we wouldn't call anything "UNIX-like" if it lacked unnamed pipes, signals or sockets.

3. **Input** – Duh.

4. **Command-line arguments** – Now we're at the subject of this column, so let's dig deeper.

A Little Background

The only Standard-C function with two alternative prototypes is `main()`, which you can use as either `int main(void)` or `int main(int argc, char *argv[])`. Whenever you use the latter, the operating system passes `main()` two parameters. The first is the number of arguments that your program was invoked with, and the second is a NULL-terminated list of those arguments.

If you're not as old as we are, you may take command-line arguments for granted. Believe us, they're a big deal. This is from Brian Kernighan's *Why Pascal is Not My Favorite Programming Language* (Bell Labs' Computing Science Technical Report No. 100):

There is no notion of access to command-line arguments, again probably reflecting Pascal's batch-processing origins. Local routines may allow it by adding non-standard procedures to the environment.

We'll give you the obligatory reminder that all arguments are processed by the

shell before calling the program. The shell does file name expansion (globbing) and splits the argument list on white space. If you say

```
$ rename foo* bar*
```

the shell expands the strings `foo*` and `bar*` before it starts `rename` (`rename` will never see the asterisks).

What do these arguments mean? Whatever you want: file names, flags, assignment statements, your children's names, your birthday and so on. The C Standard assigns them no built-in meaning.

Admittedly, unless you go to a lot of work, `argv[0]` contains the name the program was invoked by. The remainder of the arguments are the other strings from the command line that invoked the command. (Here "a lot of work" typically means invoking one of the `exec()` family by hand. If you want to know more about that, check out the man page for `exec1()`.)

This isn't much restriction, though. In the following command,

```
$ foo './abc*' -x cat=fezmo nkids=3 riley gillian zoe 11/21/48
```

the argument `nkids=3` could be a file name, as far as the C Standard is concerned:

```
$ ls -l
total 3
-rw-rw-r-- 1 jsh rd 12 Mar 29 15:08 abc
-rw-rw-r-- 1 jsh rd 24 Mar 29 15:08 nkids=3
-rw-rw-r-- 1 jsh rd 18 Mar 29 15:08 riley
```

Conventions

C isn't the only standards game in UNIX-town. Whenever we discuss command-line issues, we consult the POSIX standard for the Shell and Utilities, POSIX.2-IEEE 1003.2-1992.

In addition to standardizing commands and shell syntax, Dot 2 contains a variety of definitions and rules designed to make shell programs and programmers more portable. In particular, Section 2.10, "Utility Conventions," gives a detailed set of rules about command options. Following these guidelines, if we see this line in a shell script:

```
gut_morgn -s 'Vos makhstu?' -a -- -zoe
```

we can reasonably infer that `gut_morgn` is the name of a command, while the leading hyphens of `-a` and `-s` indicate they're flags (options) controlling the command's behavior.

'`Vos makhstu?`' is probably an argument to the `-s` option because of its position. But `-zoe` is not an option at all, despite the leading hyphen, because it follows the end-of-argument indicator `--`. Admittedly, we're talking about a standards document: precise, but not always penetrable.

Guideline 7 of Dot 2 says "option-arguments should not be optional." Fortunately, you don't have to understand it. Dot 2 also encourages conformance to its rules by providing a C-language function, `getopt()`, and a shell-level utility,

`getopts`, that encapsulate the rules of correct behavior. Use these to parse arguments, and you'll end up doing the right thing.

Of course these, too, are specified in standards-speak. Here's a paragraph taken from the four-page description of `getopts`:

If an option character not contained in the optstring operand is found where an option character is expected, the shell variable specified by name shall be set to the question-mark (?) character. In this case, if the first character in optstring is a colon (:), the shell variable OPTARG shall be set to the option character found, but no output shall be written to standard error; otherwise, the shell variable OPTARG shall be unset and a diagnostic message shall be written to standard error. This condition shall be considered to be an error

detected in the way arguments were presented to the invoking application, but shall not be an error in getopts processing.

And vice versa. Offer not good in Alaska, Albuquerque and southeastern counties of northern West Virginia between the hours of

12 p.m. and 12 a.m., Central Standard Time, inclusive. Void where prohibited by law.

Picking Arguments for Fun and Profit

The Dot 2 rationale—which contains nonnormative, explanatory and historical footnotes to the normative first volume—also provides code illustrating how to use both `getopt()` and `getopts`. (*Nonnormative* is standard-ese for "This is interesting and helpful, but language lawyers can officially ignore it." *Normative* is standard-ese for "The semicolons in the remainder of this document shall be assigned and conform to all special meanings detailed to them by section 1.3.4.2.8.7.16 of this document, except as modified by ISO Standard 1769.8a-1987, q.v., if supported by the implementation.")

Unfortunately, you probably do not have a copy of Dot 2 handy. To help smooth your way, we'll give you examples of both `getopts(1)` and `getopt(3)`.

We'll make the two code examples parallel one another as closely as we can, so you can contrast them, and we'll follow these with a third parallel example in Perl.

Let's begin with the most verbose of the three, a C program that demonstrates how to use `getopt(3)`:

```
#include <stdio.h>
#include <limits.h>
#include <getopt.h>
#include <string.h>

int r;
```

In addition to standardizing commands and shell syntax, Dot 2 contains a variety of definitions and rules designed to make shell programs and programmers more portable.


```

char g[128];

die(char *s)
{
    fprintf(stderr, "%s\n", s);
    exit(1);
}

main(int argc, char *argv[]) {
    int i, name;
    char usage[128];
    char **s;

    sprintf(usage,
        "%s: [-g value] [-r] filename [...]",
        argv[0]);

    printf("The full command line is ' ');
    for (s = argv; *s; s++) {
        printf(" %s", *s);
    }
    printf("\n");

    printf("The command name is '%s'\n", argv[0]);
    printf("There are %d other arguments",
        argc - 1);
    if (argc > 1) {
        printf(" and the first of these is '%s'",
            argv[1]);
    }
    printf("\n");

    while ((name =
        getopt(argc, argv, "g:r")) != -1) {
        switch(name) {
            case 'g': strncpy(g, optarg, 128);
                break;
            case 'r': r = 1;
                break;
            case '?': die(usage);
        }
    }

    argc -= optind - 1;
    memmove(argv + 1, argv + optind,
        (argc + 1) * sizeof(char *));
    if (argc < 2) { die(usage); }

    printf("The options are -g = '%s', -r = %d.\n",
        g, r);
    printf("The %d filename arguments are '",
        argc - 1);
    for (s = argv + 1; *s; s++) {
        printf(" %s", *s);
    }
    printf(" '\n");
}

```

Next, a shell script with a parallel `getopts` (1) example:

```

#!/bin/sh

die() {
    echo "$* 1>&2; exit 1
}

usage="$0: [-g value] [-r] filename [...]"

echo "The full command line is ' $0 $* '."
echo "The command name is '$0'"
echo -n "There are $# other arguments"
test $# -gt 0 &&
    echo -n " and the first one is '$1'."
echo

while getopts g:r name
do
    case $name in
        g) g="$OPTARG" ;;
        r) r=1 ;;
        ?) die $usage ;;
    esac
done

shift $(( $OPTIND - 1 ))
test $# -gt 0 || die $usage;

echo "The options are -g = '$g', -r = '$r'."
echo "The $# filename arguments are '$*'."

```

Finally, here's the same thing in Perl:

```

#!/usr/local/bin/perl -w
use Getopt::Std;
local ($opt_r, $opt_g) = (0, "");

$usage = "$0: [-g value] [-r] filename [...]\n";

print "The full command line is
    '$0 @ARGV '\n";
print "The command name is '$0'\n";
print "There are ", scalar @ARGV,
    " other arguments";
print @ARGV ?
    " and the first one is '$ARGV[0]'\n" :
    "\n";

getopts "g:r" and @ARGV or die $usage;

print "The options are -g =
    '$opt_g', -r = '$opt_r'\n";
print "The ", scalar @ARGV, "
    filename arguments are ' @ARGV '\n";

```

The size differences between the code are noteworthy. If we count nonblank lines, each example is about half the

size of its predecessor:

```
$ for i in getopt.c getopt.sh getopt.pl
> do
> echo ==$i
> sed '/^$/d' $i | wc -l
> done

==getopt.c
 46
==getopt.sh
 22
==getopt.pl
 11
```

Even with blank lines, the entire Perl script fits comfortably inside a 24x80 terminal window.

At the other end of the spectrum, the C program fills an entire printed page, but if you're writing a C program, you won't care how much shorter the argument parsing would be if you were writing a shell script; you will care that a `getopt()`



call is shorter, more robust and easier to debug than a hand-crafted `argv[]` parser.

But *why* are they different sizes? Comparing the C and shell sources is instructive. Almost every meaningful individual unit in the C program translates directly into something the shell can say in fewer words. For example, loops turn into single lines, variables and

functions no longer need declarations and even case statements can be laid out more directly in the shell.

Common actions have been made easy to express.

Perl is a completely different story. C's declarations and `#include` directives, which the shell script got by without, have returned. The stunning shrinkage comes from something else: much-simpler `getopts` handling. Parsing the options, storing their values, verifying that there are file name arguments and issuing a usage message for improper invocation takes 13 lines in C, 11 lines in the shell and only one in Perl.

If we stand back and remind ourselves that the bulk of these three programs are just scaffolding to report that we've parsed the arguments successfully, the author of `Getopt::Std` has clearly done something right.

More Arguments

But what if you need to do really, *really*, *really* complicated argument processing?

In the shell, you're back to doing things by hand. In C, take a look at the man page for GNU's `getopt_long()`, which can do so many things that the Bugs section reads:

This man page is confusing.

Perl, too, has a comprehensive, confusing, `Getopt::Long` module, but if you don't like it, the Comprehensive Perl Archive Network (CPAN) at <http://www.perl.com/> lists seven other packages, which offer even more styles of argument handling. However, even if you choose to use `Getopt::Std`, get the latest version from the CPAN. Versions before 5.004 only let you save the value of the `-d` option into the variable `$opt_d`. Recent versions let you save arguments into the hash (associative array) of your choice. Instead of saying this:

```
getopts('abcdefg') or die $usage;
print "\$opt_a is $opt_a\n" if defined $opt_a;
print "\$opt_b is $opt_b\n" if defined $opt_b;
...
```

you can now say this:

```
getopts('abcdefg', \%opts) or die $usage;
foreach (sort keys %opts) {
    print "\opt{'$_'} is $opts{$_}\n"
}
```

Does it matter what hash we choose? There is one interesting choice: the hash `%ENV`, which contains the environment. By using `%ENV`, you can have your code set options from the command line, *or* take in option values from the environment. Observe:

```
# Set the default values
# override environment value
$ENV{'r'} = 0;
# set to "hello" only if $g
# isn't already set in the environment
$ENV{'g'} ||= "hello";

# optionally override defaults
# with command-line arguments
getopts($args, \%ENV)
or die "usage: $0 [-r] [-g string]
```

Of course, this takes us to the beginning of our column. We'll let you reread it while we go off and argue until next month.

Until then, happy trails. →

Jeffrey Copeland (copeland@alumni.caltech.edu) lives in Boulder, CO, and works at Softway Systems Inc. on UNIX internationalization. He spends his spare time rearing children, raising cats and being a thorn in the side of his local school board.

Jeffrey S. Haemer (jsh@usenix.org) works at QMS Inc. in Boulder, CO, building laser printer firmware. Before he worked for QMS, he operated his own consulting firm, and did a lot of other things, like everyone else in the software industry.

Note: The software from this and past Work columns is available at <http://alumni.caltech.edu/~copeland/work.html>.

Graphics have come to corporate networks, thanks to large investments in backbone bandwidth, but most company network infrastructures still won't support the next frontier—video applications.




CHRIS BUTLER

INTRANETS STILL NOT READY FOR VIDEO

by PAUL KORZENIOWSKI

Corporations have invested billions of dollars during the past few years to build up their network infrastructures. A year or two ago, 50 to 100 users shared one 10-Mb/s Ethernet connection, and users dumped interdepartment data onto backbone connections with a total of 100 Mb/s of bandwidth. The advent of Ethernet switching pushed 10-Mb/s bandwidth down to many users' desktops and made 100-Mb/s backbone connections quite common—in fact, a growing number of corporations are now dabbling with 1-Gb/s backbones.

This influx of low-cost, high-speed bandwidth has enabled companies to install high-resolution graphics and imaging applications. Users now regularly work with Internet and intranet systems that present complex graphics with a simple click or two of the mouse.



Yet there is one application that network infrastructures still cannot easily support: video. Video systems require gobs of bandwidth and special networking features that are not yet widely available. While there has been plenty of hype about the benefits video offers, corporations have shown little interest in the technology and it is unclear if or when that will change. "Currently, most companies are just testing video in research and development groups," says Ron Lifton, president of Class Data Systems Inc., a Cupertino, CA-based network software supplier.

Deployment problems start with the amount of bandwidth these applications require: a typical 30-second video clip can take 25 MB of data. While corporate networks can now handle bigger workloads than ever before, few have sufficient bandwidth to move so much data with ease. Because video is such a bandwidth hog, corporations have taken steps to ensure it will not affect other applications. In most cases, video servers have been placed on autonomous

networks. Building separate networks is expensive, so companies have limited the number of video connections and placed videoconferencing systems in common areas, such as conference rooms.

But for video to reach its potential, video applications must be available to every user in an organization. They have to make their way onto users' desktops. For this to occur, vendors must address several technical issues.

Quality First

First, the quality of video applications must improve. Ed Buckingham, analyst at International Data Corp.'s Mountain View, CA-based, office, says, "Earlier this year, I watched Intel [Corp.]'s Andrew Grove conduct a videoconference, and frankly, the picture quality was not good: It was jumpy and had poor resolution."

But quality is a tricky problem for vendors to solve. Video applications transmit data in frames per second; the more frames sent, the clearer a picture becomes. The downside is the more

frames sent, the more bandwidth needed.

Desktop video systems now transmit information at rates of 20 to 22 frames per second (fps), which take up 384 to 512 Kb/s of bandwidth. Suppliers want companies to accept lower transmission rates: 12 to 15 fps, which would mean 128 Kb/s of bandwidth. But at this rate, a picture may appear jumbled if a user moves around quite a bit. (By comparison, a standard television in the United States operates at approximately 30 fps.) Even if companies are willing to accept a reduction in quality—and there is no assurance they will—video applications represent a quantum leap in network bandwidth use: ten times that required to paint a Web page.

So users will not be able to run desktop videoconferencing applications over shared Ethernet networks where groups of users work with one 10-Mb/s network. Instead, a company has to upgrade its connection to at least switched Ethernet, where each user has a dedicated 10-Mb/s connection—this will provide good but not great video connections. For quality video, corporations may have to deploy 100-Mb/s desktop connections.

Adding bandwidth addresses is one issue, but a second set of problems stem from the nature of videoconferencing applications. Currently, information travels across Ethernet networks in random fashion. In a series of 10 packets, packet number eight may arrive before packet number six and so on. The computer system at the other end shuffles the packets so the information is presented correctly to an end user.

Packet arrival order is not important with most data applications but is vital with video transmissions. If packets arrive in the improper sequence, a picture may fluctuate or images may appear jumbled. If two workers are conversing, they may find the fluctuation annoying. If two doctors are examining medical images, for example, the flickering could be catastrophic.

A related problem is bandwidth contention. On an Ethernet network, bandwidth is parceled out on the fly. A user may begin sending a large file when no one else is using the network and the transmission will start out fine.

Freebie

This is a free **CD OFFER** for people interested in software and hardware products for the Sun environment. The **SUN SOLUTIONS CD** integrates the power of CD-ROM technology and seamless interaction with the World Wide Web utilizing **JAVA™**. Inside each issue, you'll discover a wide range of product and industry information, a complete catalog of all available software for the Sun technologies, product demos... even free software in most issues. Sign up for a **FREE SUBSCRIPTION** to this comprehensive information tool.

Subscribe Today:

email: sunsolcd-form@sun.com

<http://www.sun.com/sunsolcd/>

fax back: +1-510-372-8582



Circle No. 27



Video and Networking

A colleague may start to access a database, and the transmission could slow to a crawl. With a file transfer, the only impact is that users have to sit and wait a little longer than they would like.

Video applications cannot tolerate such fluctuations. If two users are conversing and the available bandwidth shrinks, a transmission will jar or possibly break completely. Video applications need to have bandwidth available during an entire session.

A networking feature called quality of service (QOS) solves some of these problems by opening up a clear communication line between two end points so data can move freely. QOS is an emerging network feature available on select products. This capability was first delivered back in 1995 on Asynchronous Transfer Mode (ATM) networks, and supporters thought it would convince companies to deploy ATM at users' desktops. However, such a transformation has not occurred. Corporations have opted to outfit their desktops with Ethernet, which, according to Dell'Oro Group, a Portola Valley, CA-based research firm, accounted for 83% of the revenue network equipment suppliers generated in 1997 (see Table 1).

For many companies, Ethernet promises less disruption than ATM. During the past decade, organizations have purchased management applications and trained scores of network technicians to monitor Ethernet networks. To deploy ATM, a company has to scrap those investments, install new adapter cards, wire hub connections and network management software and train its technicians to oversee a second type of network.

ATM also costs more than Ethernet. The price for a 10/100-Mb/s Ethernet adapter has dropped to less than \$200 (vendors such as Sun Microsystems Inc. include them with their computers for free) and a terminal-to-hub connection is priced at less than \$300. In comparison, a 155-Mb/s ATM network interface card costs \$500 to \$1,000, and connection prices hover between \$750 and \$1,500. ATM supporters developed a 25-Mb/s option, which brought the

adapter card price down to the \$300 range and the connection charge to less than \$1,000, but corporations have shown little interest in this option because Ethernet offers more bandwidth at a lower price.

Fore Systems Inc., Warrendale, PA, still thinks ATM will gain widespread desktop acceptance. Stephen Duffy, a product manager for video and multimedia applications at the network equipment supplier, says, "Corporations are starting to realize that running video on ATM networks eliminates problems such as bandwidth contention, fluctuating pictures and packets being dropped."

But most observers think Ethernet has already won the desktop battle. The first products from First Virtual Corp., a Santa Clara, CA-based start-up that builds software to help corporations deploy desktop videoconferencing applications, ran on 25-Mb/s ATM connections. In 1997, the company shifted its development efforts to Ethernet. "Any vendor's strategy that is based on the

belief that users are going to migrate from Ethernet to ATM at the desktop will fail," says Jim Nielsen, vice president of marketing at First Virtual.

The company was able to change its plans because an Ethernet alternative to QOS is now under development. The Internet Engineering Task Force (IETF), an ad-hoc standards group that has guided the Internet's evolution, began working on Resource Reservation Protocol (aka RSVP), which adds QOS service to TCP/IP—and therefore Ethernet—networks.

This option is gaining a great deal of support. Networking equipment suppliers Bay Networks Inc., Santa Clara, CA, Cisco Systems Inc., San Jose, CA, and 3Com Corp., Santa Clara, CA, plan to add RSVP features to their networking equipment by the end of this year.

How well RSVP will work is debatable. Because TCP/IP was not initially designed to support QOS, the emerging feature could add a lot of network overhead and decrease network performance.



Concorde Group, Ltd. Announces the U-MAX5 and U-MAX10 Product Series

What makes Concorde Group's U-MAX product series the right choice?

Sun Microelectronic's™ revolutionary Ultra AXI™ runs a pure native version of Solaris™ embedded into a PCI form-factor. The U-MAX5 and U-MAX10 product series are Concorde's latest UltraSPARC™ PCI based offering. The price/performance ratios of these products are setting the standard within the industry. See for yourself!

- Modular CPU Design for easy inexpensive upgrades to prolong your company's hardware investment
- 6 PCI slots provide maximum expandability and bandwidth for I/O intensive networking applications
- Integrated Dual-Channel Ultra-Wide SCSI for increased data transfer rates up to 40MB/sec
- Advanced System Monitoring (ASM) automatically detects system problems before they affect the system operation
- Graphic Options range from standard SVGA to Creator 3D and FFB2 for the most graphic intensive applications
- PC Tower, Desk side Tower and 19" Rack Mount with removable disks and internal storage capacities up to 180GB

Concorde U-MAX products are field serviced by CV Support International™ and Polaris Services, Inc.™

This is NOT a compatible!
This is NOT a clone!



This is a 100% Sun™ manufactured motherboard and processor



CONCORDE GROUP, LTD.

753 Forest Street, Marlborough, MA 01752
Phone: (508) 804-0400, Fax: (508) 804-0500



ALL TRADEMARKS MENTIONED ARE PROPERTY OF THEIR RESPECTIVE COMPANIES.

For more information: 1-800-333-2786 <http://www.concorde.com>

Circle No. 28

Video and Networking

Fore Systems' Duffy adds that RSVP lacks needed features. For instance, the protocol can initially allocate bandwidth to an application but cannot ensure it will continue to be available as transmission occurs—ATM networks can perform both functions.

The debate between RSVP and ATM QOS hinges on the belief that many corporations will deploy QOS functions. "I would estimate that less than 5% of our customers have added QOS capabilities to their networks, but I expect the percentage to grow significantly in the future," says Arvie Martin, a senior product manager for digital media at Sun in Mountain View, CA.

Yet that prediction may not turn out to be accurate. Deploying QOS is difficult. A company has to examine all its applications, determine which are most important, and configure its switches to recognize and give priority to certain applications. Network technicians must perform this task on each switch, a daunting task in medium and large-size companies.

Aware of the difficulty, network equipment vendors have been developing new management tools, called policy-based systems, that automate the process. Cisco Systems and 3Com plan to deliver such products during

the second half of this year. If QOS networking features become easier to deploy, corporations will still need one more building block for their video applications. Desktop applications need to be able to take advantage of the feature. To date, this has been a problem because desktop operating systems do not include QOS interfaces. Consequently, application developers have to add a middleware layer to their products that maps them to these features. According to industry analysts, most operating system vendors plan to add QOS capabilities to their products during 1998. Another plus is vendors have been rallying around standard interfaces for video applications (see "Desktop Video Standards").

Because there are so many missing pieces, video has been making little headway in most organizations. IDC's Buckingham says revenue from desktop videoconferencing systems was flat in 1997 and expects that to continue to be the case in 1998. But video technology has made progress in a few niches—distance learning enables students in remote locations to take special classes, corporations use video in interactive training classes and doctors in different locations can collaborate on patient care with telemedicine applications.


The College of Oceanic and Atmospheric Science at Oregon State University in Corvallis is one such video user. Chuck Sears, manager of research computing, says the college wanted 300 users to be able to work with video depictions of weather information rather than manually sift through reams of information stored in database management systems.

Because most of the users work with UNIX workstations, the college selected Sun's ShowMe video application. Sears says the group collaboration package lets users either manipulate video information in real time or download archived data.

Before deploying the application, the organization overhauled its network infrastructure. The university had relied on shared 10-Mb/s Ethernet desktop connections and upgraded to shared 100-Mb/s connections. Yet Oregon State University had no plans to add QOS features to its network.

"The video transmissions take up 4 Mb/s to 5 Mb/s of bandwidth, which we think our 100-Mb/s connections should be able to handle," says Sears. At this writing, half a dozen users were testing the application, and the university expected to roll it out during the summer.

Table 1. Ethernet Continues to Dominate U.S. Networking Market
(in millions of dollars)



Most industry observers think Ethernet has already won the desktop battle and expect that to continue to be the case into the next millennium.

Source: Dell'Ono Group, Portola Valley, CA

Video and Networking

The university can justify the major network upgrade needed to support the video applications, but most corporations cannot. "To gain wider acceptance, video applications have to provide corporations with clear-cut business benefits," says IDC's Buckingham. Vendors have focused on the savings from reduced travel that desktop video-conferencing makes possible, but that has not been a very convincing argument.

Video via Email

Steve Fingerett, product marketing manager at the Minneapolis, MN-based office of Warner Brothers Records Inc., has found a strong business reason for deploying video applications. A self-proclaimed technojunkie, he is constantly on the lookout for tools that simplify his job, which is to convince radio station programming managers to play music from Warner Brothers Records' artists.

In February, a friend told Fingerett about CVideo-Mail from Cubic Video-Comm Inc., a subsidiary of Cubic Corp., San Diego, CA. The product enables users to insert video clips in electronic mail messages and includes an embedded video player so recipients can watch the videos.

In March, Fingerett inserted portions of Madonna's "Frozen" music video into a message and sent it off to a handful of programming managers. "Some of these managers did not get to see the video when it debuted on MTV, so my electronic mail message enabled them to stay current with what listeners had seen," Fingerett says.

Inserting the clip was a simple cut-and-paste exercise. Fingerett says one limitation is users are unable to edit the clips, a limitation he expects Cubic to address in a future release of the product. Fingerett is now experimenting with the product and expects to use it more in the future. There is even a possibility it will be deployed companywide.

Eventually, proponents expect more companies to identify applications that can benefit from video transmissions. "Because timely access to information has become a top concern in just about every company, managers are interested in applications such as downloading

Desktop Video Standards

Industry standards make it simpler for corporations to connect different vendors' equipment. In the last few years, video suppliers have worked on such specifications. The International Telecommunications Union has taken charge of outlining standards for desktop systems running on telephone lines, ISDN connections and LANs. Here is a brief description of the key standards and their current status.

◆ **T.120:** A specification that allows desktop users to hold a conference spanning multiple locations. It enables users to view files and move them from one location to a second via FTP. Support for this specification is included in many current desktop video systems.

◆ **H.320:** This standard enables desktop videoconferencing audio and video components to be shared in multiple locations over ISDN telephone lines. Products complying with this specification began arriving in 1996.

◆ **H.323:** A specification for transmitting videoconferencing sessions on

LANs, such as Ethernet. Products conforming to this specification are expected to arrive this year.

◆ **H.263 and G.723:** These two compression standards are subsets of the H.323 specification. G.723 works with audio transmissions and H.263 compresses a 30-MB video file (containing roughly 30 seconds of video) to a 100-Kb/s transmission.

◆ **H.324:** This standard allows users to hold videoconferences on analog telephone lines via modems. Because the top transmission speed of 28.8-Kb/s provides low-quality video transmissions, interest in this standard and compliant products has been limited.

◆ **MPEG1:** A format so desktop applications can recognize and work with video files. Support for this standard is common in group collaboration packages.

◆ **MPEG2:** A follow-on to MPEG1 that promises greater picture clarity and support for interactive transmissions. Products that support this standard should start to arrive this year.—pk

"Stay ONLINE with INLINE"



Your Key to

RAID

Why do Fortune 1000 Corporations, VARS, and System Integrators use INLINE.....



800-465-4637

<http://www.INLINECORP.com>

- 100 Plus Megabytes per Second
- 10,000 RPM Drives
- 12 Gigabytes to over 8 Terabytes
- Ultra SCSI & Fibre Channel
- Windows NT, UNIX, Netware & Others
- Multiple Host & Network Attach
- RAID Levels 0, 1, 3, 4 & 5

INLINE Corporation • 585 Grove Street Suite 310 • Herndon, Virginia • 20171 • Main 703-478-0800 Fax 703-478-0800

SunExpert Magazine ■ June 1998

Circle No. 29

daily financial news broadcasts to employees' desktops," says Sun's Martin.

Because there are so many hurdles to overcome in deploying these applications, two to five years may pass before they are widely implemented. "Video acceptance has occurred much more slowly than we anticipated, yet there has been progress," says First Virtual's Nielsen. "The market's evolution reminds me of the talk about LANs in the late 1980s. Every year, there were predictions that this year would be The Year of the LAN. There was no one Year of the LAN, but by the early 1990s, most corporations had installed them. We don't anticipate one year when corporations will jump on video, but we expect that in a few years most organizations will be running video applications." ➔

Paul Korzeniowski is a freelance writer based in Sudbury, MA, who specializes in networking issues. His email address is paulkorzen@aol.com.

Companies Mentioned in this Article

3Com Corp.
5400 Bayfront Plaza
Santa Clara, CA 95052
<http://www.3com.com>
Circle 150

Bay Networks Inc.
4401 Great American Pkwy.
Santa Clara, CA 95052
<http://www.baynetworks.com>
Circle 151

Cisco Systems Inc.
170 W. Tasman Drive
San Jose, CA 95134
<http://www.cisco.com>
Circle 152


Class Data Systems Inc.
19200 Stevens Creek Blvd.
Cupertino, CA 95014
<http://www.classdata.com>
Circle 153

Cubic VideoComm Inc.
9333 Balboa Ave.
San Diego, CA 92123
<http://www.cubic.com>
Circle 154

First Virtual Corp.
3393 Octavius Drive, Ste.102
Santa Clara, CA 95054
<http://www.fvc.com>
Circle 155

Fore Systems Inc.
1000 Fore Drive
Warrendale, PA 15086
<http://www.fore.com>
Circle 156

Sun Microsystems Inc.
2550 Garcia Ave.
Mountain View, CA 94043
<http://www.sun.com>
Circle 157



READER FEEDBACK

To help *SunExpert* serve you better, take a few minutes to close the feedback loop by circling the appropriate numbers on the Reader Service Card located elsewhere in this magazine. Rate the following column and feature topics in this issue.

Features:

Intranets Still Not Ready for Video
Review: Cross-Platform Storage for the Masses

Columns:

UNIX Basics—make: Mastermind of the Update
NTegration—Sharing File Systems, Part 2
Datagrams—Processing Masses of Mail
Work—We Argue with Our Code

WebServer Magazine:

UNIX vs. NT Myth
Communicating without a Net
Clustering
A Good Web Server Platform at the Right Price

Interest Level

High	Medium	Low
175.....	176.....	177
178.....	179.....	180
181.....	182.....	183
184.....	185.....	186
187.....	188.....	189
190.....	191.....	192
193.....	194.....	195
196.....	197.....	198
199.....	200.....	201
202.....	203.....	204

Escape Beeper Burnout!!!

Ensure Application Availability with PATROL®



Are 9-1-1 calls about crashes and slow-downs with your mission-critical applications spoiling your dinner? Don't let it happen. The proven combination of PATROL enterprise management products from BMC Software, and the power of Sun Microsystems will let you enjoy your meal, and your life, by guaranteeing system availability, performance and recovery across your enterprise.

Your system is more than the application; It's the operating system, messaging and middleware, databases and more. PATROL plugs into all of these and provides out-of-the-box solutions that can free you from beeper burnout. Maximize Sun Solaris' scalability and reliability with BMC Software's cost effective PATROL solutions and get application service assurance 24 hours a day, 365 days a year.

Whatever your system environment, PATROL solutions are ready to start simplifying your life right now. Call the Sun registration hotline at 800 622-4786, mention PATROL and find out how to make your life more manageable while saying goodbye to beeper burnout.

Take the time to answer a few simple questions when you call and we'll send you a PATROL Interactive CD and a *cool* T-Shirt – **YOURS FREE** just for calling.



BMC Software is an Equal Opportunity Employer. BMC Software, the BMC Software logos, and all other product or service names are registered trademarks or trademarks of BMC Software, Inc. in the USA and in other select countries. © and ™ indicate USA registration or USA trademark. Sun, Sun Microsystems and Solaris are registered trademarks or trademarks of Sun Microsystems. Other logos and product or service names are registered trademarks or trademarks of their respective companies. © 1998 BMC Software, Inc. All rights reserved.



Cross-Platform Storage for the Masses

by IAN WESTMACOTT, Technical Editor

This month, we put TAS 5.2—a storage solution that is winning praise from both users and vendors—through its paces and, in the final analysis, it's definitely worth a look.

When it comes to sharing storage interactively across popular platforms, solutions abound. A plethora of application and protocol software solutions, as well as dedicated hardware solutions, give system planners and managers many choices, but little in the way of absolute or even relative performance and feature comparisons. Syntax Inc.'s TotalNET Advanced Server (TAS) 5.2 is a solution that is winning praise from users and vendors alike.

TAS enables UNIX servers (AIX 4.1.4, 4.1.5 and 4.2; Solaris 2.5 and 2.6, SPARC or i386; HP-UX 10.10 and 10.20; IRIX 5.3 and 6.x) to become file, print and application servers for PC clients via native file system and network transport protocols. Client platform support includes Macintosh, MS-DOS, Windows for Workgroups, Windows 95/NT and OS/2 Warp. TAS implements three network transport "realms": LAN Manager (LM)-NT-OS/2 for NetBIOS clients over TCP/IP or NetBEUI, NetWare for NetWare clients over IPX/SPX, and AppleTalk for Macintosh clients over AppleTalk. TAS also supports the Common Internet File System (CIFS) and Distributed Computing Environment (DCE).

New in Version 5.2 is the TotalNET Administration Suite (TNAS) and TotalAdmin, a revamping of the Web browser-based administration of TotalNET; support for case-sensitive and long file names, alleviating difficulties with multiplatform applications and TAS utilities; improved TAS utilities' performance; enhanced tracing options for monitoring and debugging; AppleTalk printer access from non-AppleTalk realms; and online documentation in HTML format. Also new is an August 1997 agreement with Sun Microsystems Inc. to include a single-user copy of TAS with Solaris for Intranets, and a January 1998 agreement with IBM Corp. to include a single-user copy of TAS with AIX. TAS is also bundled with some Silicon Graphics Inc. systems. Customers may purchase additional licenses from Syntax.

Transparent Client Access

One of the most attractive features of TAS is transparent client access. No modification of client systems is necessary; no client-side configuration or software is required. Once the server is installed and configured, clients can browse TAS servers and services immediately using the client's native network and file system protocols. For example, users can browse file and printer services via the Chooser in Mac OS or via the

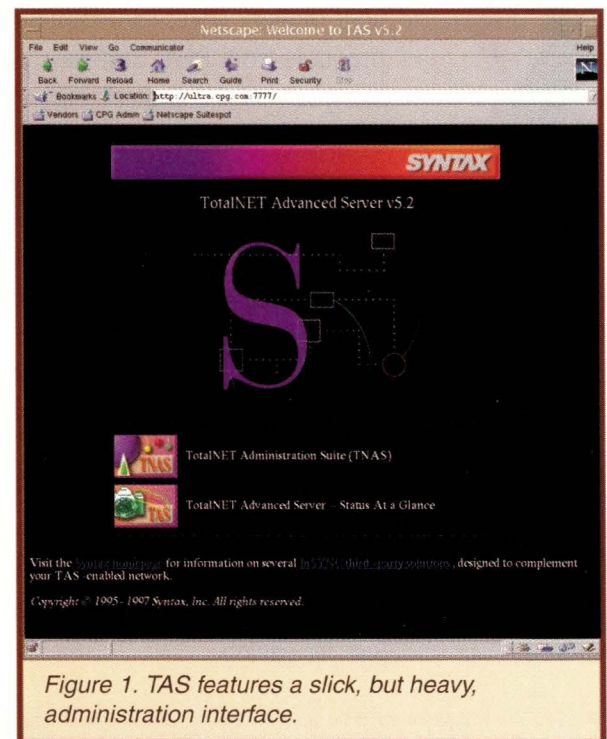


Figure 1. TAS features a slick, but heavy, administration interface.

Network Neighborhood for Windows 95/NT clients. The server accomplishes this by loading the appropriate network transports for each realm implemented and by mapping client file system attributes onto the UNIX file systems, which store the shared resources.

Because client file systems in the various realms implement different file attributes, TAS makes use of “shadow attribute” files and directories to map these attributes onto the UNIX file system. For example, the Mac OS file system stores two “forks” for data files: a data fork and a resource fork. The data fork contains the file’s data and the resource fork contains file attributes such as type and application. TAS stores the data fork on the server as a regular file and stores the resource fork in a shadow attribute subdirectory. This mapping is completely transparent to clients within the same realm, but care must be taken when accessing shared resources through foreign realms.

The TAS server uses the Server Message Block (SMB) protocol to communicate with PC clients in the LM-NT-OS/2 realm. The SMB specification provides for interoperability between PCs and X/Open-compliant systems, as defined by The Open Group. Using this protocol, TAS is able to provide services transparently to PC clients. However, TAS must rely on proper client-side SMB implementation. Care must be taken when accessing files from other client types and realms: Different clients may implement SMB differently and support different file attributes.

For example, a Windows NT client that accesses a file saved by a Macintosh client may or may not correctly maintain the resource fork shadow attribute of that file. A file move or rename will correctly maintain the shadow attributes, but a file copy will not. According to Syntax, Windows NT implements a file copy as open, read, create and write. When the copy is created, the TAS server has no way of knowing

whether the NT client is creating a new file or a copy of an existing file. Thus, the server cannot correctly maintain the shadow attributes of the file.

On the other hand, the OS/2 LAN Requester copies files by using the SMB-copy operation, rather than reading the old file and creating a new file. Thus, the TAS server knows that a copy is being made and can correctly maintain the shadow attributes of the file. If you access the file from a UNIX client, either by UFS or NFS, then all bets are off

fairly detailed knowledge of networking and file system protocols.

A large suite of command-line utilities are provided for administering the server, but most interactive administration will be done through the Web browser-based TNAS. TNAS is a general framework for administering local and remote Syntax and third-party services, including TAS. Use of this interface requires a browser that supports Java, JavaScript, frames and HTML cookies, such as Netscape Communications Corp.’s Navigator or Microsoft Corp.’s Internet Explorer.

Although the layout and organization of TNAS is an improvement over previous versions, in my opinion there is gratuitous use of animated GIFs and applets in the interface. For example, the entrance screen has a 360-KB “flying S” animated GIF (see Figure 1), and subsequent pages incorporate Java applets whose sole function is to change color when the mouse passes over them. Design issues aside, the TNAS interface is fairly

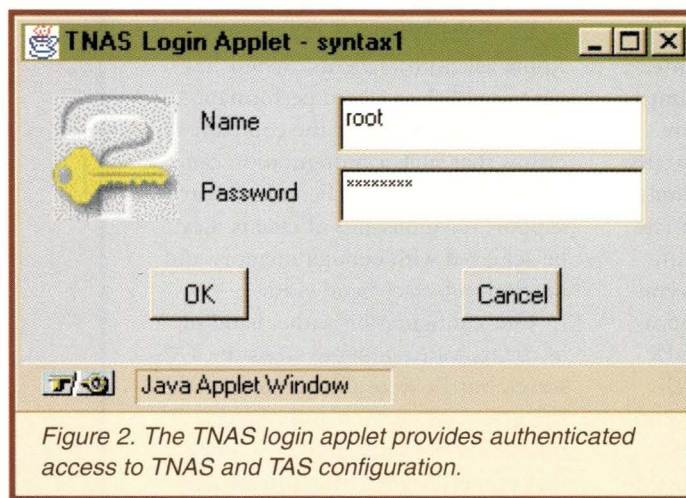


Figure 2. The TNAS login applet provides authenticated access to TNAS and TAS configuration.

because the TAS server doesn’t even know that you’re accessing the file, let alone what you’re doing with it. However, there are UNIX command-line utility programs for converting files between various client file system formats.

Administration and Documentation

Installation of TAS is typical. On Solaris systems, TAS is provided as a software package, which is installed using the `pkgadd` utility. The installation requires approximately 35 MB of disk space. Licensing of TAS is hostID-based. Apart from release and installation notes and a small quick-reference guide, all documentation is provided in HTML format. TAS comes with an HTTP server, which by default runs on port 7777 and is used to access online documentation and administration. The online documentation is rather rough, laid out like a reference manual, and not searchable. There are a number of cases where undefined terms are used and the documentation assumes a

sophisticated.

Once connected to the TNAS interface, you have two choices. The “Status at a Glance” option provides unauthenticated access to the status of the servers, indicating which realms and services are up and how many users are connected. The “Administration Suite” option provides authenticated access to TNAS configuration, as well as TAS and any other third-party product configuration. To use this option, a user must authenticate himself by way of a Java applet (see Figure 2), which stores a unique cookie in the client browser. This cookie has a configurable inactivity time-out (default 20 minutes), after which the client will be forced to reauthenticate.

TNAS maintains a notion of state for each session connection. Thus, if your session expires, when you reauthenticate you can go back to the same point where you left off. There is also a configurable session time-out (default 480 minutes), which is the amount of inactive time TNAS will maintain the session state. This model requires you to use TNAS

Product Review

navigation buttons, rather than browser navigation buttons, in order to correctly maintain session state. Also, an explicit logout button is provided to force the session to expire when you are finished with configuration.

TAS configuration within the TNAS framework is similar to previous versions in which a hierarchical set of configuration screens for the system, realms, transports and users is available. Most of these pages use simple forms-based user input. In addition, hyperlinks to the documentation are provided for most options, although in numerous instances the links pointed to the wrong sections of the documentation. Three permissions levels are provided: root user, TAS administrator and generic user. TNAS displays only the configuration options appropriate to the permission level of the client.

There are two choices for access control of TAS services: UNIX authentication and secure authentication. UNIX authentication uses standard UNIX mechanisms, mapping client users onto the `passwd` database (could be NIS or some other directory service). Using this method, passwords are sent over the wire as clear text. User maps may be used to map a group of users onto a single UNIX account. In the secure authentication scheme, the TAS server maintains its own database of users, and passwords are encrypted over the network.

Users of TAS 5.x are entitled to a free 10-user version of TunTAS, a server-based terminal emulation package for Windows clients, providing vt320 and wyse60 terminal support in a number of languages. TunTAS is based on Esker Inc.'s Tun PLUS product line and

can be upgraded to support additional emulation types and functionality. Also available is TAS-DCE, a module to allow OSF DCE users to log into TAS and share resources from DCE DFS. TAS-DCE uses Distributed Security Services (DSS) for access control and DFS to make TAS services available to DFS clients.

Performance

We installed TAS on a Sun Ultra 1 170 with 128 MB of RAM and used the LM-NT-OS/2 and AppleTalk realms from Windows 95/NT and Apple Macintosh clients. Syntax has yet to publish any hard performance numbers for TAS, but the company claims that such a configuration can support hundreds of clients and that support for thousands of clients may be achieved with enough memory and network interface bandwidth.

Our configuration with a handful of clients wasn't going to stress the TAS server, but we were able to measure the resource consumption and response time of the TAS daemons. We found that the server grew by approximately .3 MB per connection under moderate use, validating Syntax's claims. We also found that response time was very good, approximately 30% faster from a Macintosh client than a 200-MHz Pentium Pro running Windows NT and Services for Macintosh.

Whether you wish to share storage among different types of PC clients or provide centralized storage for a large number of PC clients, TAS definitely deserves a look. Almost plug-and-play, all configuration is performed on the server, providing seamless access to

Windows, NetWare and Macintosh clients. While the interface could lose some of the extraneous frills, the Web-based management and configuration is one of the best we've seen. And for those of you who have a command-line bent, or want to automate management functions, command-line utilities are provided for most management functions. ➡

TotalNET Advanced Server 5.2

Company

Syntax Inc.
840 South 333rd St.
Federal Way, WA 98003

Phone

(253) 838-2626

Fax

(253) 838-9836

URL

<http://www.syntax.com>

Price

Server/initial connection	\$0
Five-user license	\$849
25-user license	\$2,399
50-user license	\$4,499
100-user license	\$8,499
500-user license	\$27,999

Best Feature

Transparent heterogeneous client access to services

Worst Feature

Documentation

Circle 158

BEGINNING THIS MONTH

WEBSERVER MAGAZINE SUPPLEMENT

Taken from our *WebServer OnLine Magazine*, the monthly editorial section will be directed to those readers who are involved with evaluating hardware/software products for corporate Web sites and Internet applications.

WebServer

For Managers of World Wide Web Sites

M A G A Z I N E

68 *UNIX vs. NT Myth*

Ian Westmacott, Technical Editor

Ever since its launch, Windows NT has been held up against the venerable tried-and-true UNIX. Both sides have avid proponents and the popular press has been awash in the UNIX vs. NT debate.

73 *Communicating without a Net*

Simson L. Garfinkel, Contributing Editor

The promise of wireless data communications is simple: free bandwidth. Just buy a pair of wireless Ethernet bridges, put antennas on a couple of roofs, and you'll have the equivalent of a T1 data circuit between two buildings without spending another dime.

78 *Clustering*


Ian Westmacott, Technical Editor

No matter how big a server is, no matter how fast the processor or how much memory is installed, there are still numerous bottlenecks that affect performance and represent single points of failure. Clustering techniques alleviate these problems for high availability and load balancing of service.


81 *A Good Web Server Platform at the Right Price*

Ian Darwin


Looking for a Web server platform that is affordable, reliable, secure, scalable and has a list of satisfied users? Then look no further than OpenBSD.



Page 68



Page 78




Page 81

Visit our Web site at <http://webserver.cpg.com>
for archived articles and a listing of new Web products.

The UNIX vs. NT Myth

Ever since its launch several years ago, Windows NT, Microsoft Corp.'s New Technology operating system, has been held up against the venerable tried-and-true UNIX. Both sides have their avid proponents, and the popular press has been awash in the UNIX vs. NT debate.



Now, with the immense popularity of the World Wide Web and TCP/IP client/server computing, the UNIX vs. NT debate has been rekindled as IS managers search for the ultimate Web server operating system. A Web server platform, the complete hardware and software system on which a Web server is built, is a complex system comprising numerous components, each of which has its own performance issues. The operating system is the glue that fuses all these components and is itself a critical component of any Web server. The choice of operating system is perhaps the most important decision you can make, and the two names leading the pack are UNIX and Windows NT. But to make that decision, you might be best served by looking inside your own organization.

NT: State of the Art

Windows NT is a relatively young operating system that was able to take advantage of lessons learned and recent advances in operating system design. Borrowing from, among others, VMS and Mach, NT is an ambitious effort, including state-of-the-art research. Much of the NT design team trod the halls of Digital Equipment Corp., including team leader David Cutler, whose résumé includes VMS and RSX-11M. The challenge they faced was to design a next-generation operating system, conformant to a range of stringent constraints, that was as software and hardware backward-compatible along the Intel Corp. line as the Intel chips themselves. This included the requirement that Windows NT provide Win32 (NT's primary programming interface), MS-DOS, 16-bit

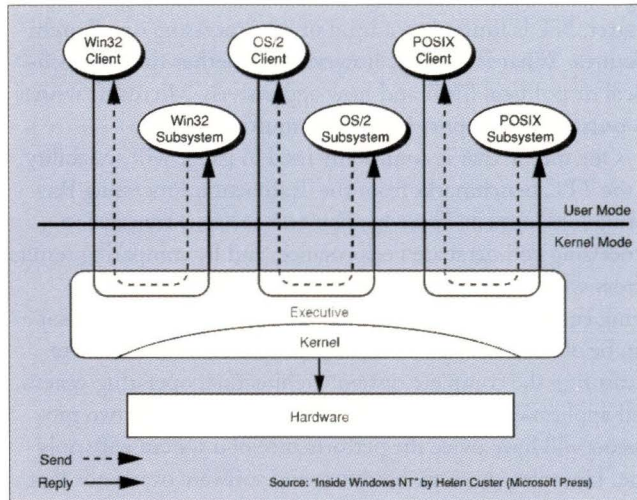
Windows, POSIX (1003.1) and OS/2 APIs, along with all the OS support that these require. Not a simple task.

The solution they chose was to use a variant of the Mach microkernel client/server design. In this approach, many of the typical operating system tasks, such as the file system and memory manager, are separated from the operating system proper, or kernel. These components then run as stand-alone server processes, fulfilling requests of other server processes or user programs. This design provides portability, ease of maintenance and support for distributed architectures. For example, various server components could run on separate nodes in a network.

Mach grew out of the Accent operating system developed at Carnegie Mellon University and was originally developed inside the BSD UNIX kernel. Besides NT, Mach was the basis of the Open Software Foundation's OSF/1 and NeXT/OS, and is the basis of the GNU HURD. Mach was designed with explicit thought given to multiprocessor systems (unlike UNIX), distributed operation and multiple operating environments.

Windows NT is not a microkernel design, but it does make use of the client/server model. NT is roughly divided into two parts, NT executive, the kernel portion of the operating system based on a layered design, and a set of "protected subsystem" servers that run in user mode. These servers each provide an operating system environment to user applications (see Figure 1). The advantages of such a design include the ability to upgrade OS environments separately from the operating system itself, the ease of adding new environments, increased security and reliability and selective porting of environments.

Figure 1. Server Mode



NT is roughly divided into two parts: the Kernel and a set of 'protected subsystem' servers that run in user mode.

Another design goal of NT was portability. Written primarily in ANSI C, with platform dependencies isolated in a hardware abstraction layer, NT was meant to easily port to multiple architectures. As evidence of its success, NT has been successfully ported to the DEC Alpha, MIPS and PowerPC architectures. But, as evidence of its commercial roots, official support for the PowerPC and MIPS processors will be dropped in Version 5 of the operating system due to lack of demand, according to Microsoft.

UNIX: An Old Friend with Many Faces

UNIX is a comparatively old operating system, initially coming out of AT&T in the late '60s and going through various states of commercialization and numerous owners—and even several main threads of development. Today, there are literally dozens of UNIX variants, including Sun Microsystems Inc.'s Solaris, IBM Corp.'s AIX, Silicon Graphics Inc.'s IRIX, Hewlett-Packard Co.'s HP-UX, DEC's Digital UNIX, Linus Torvalds' Linux, Andrew Tanenbaum's Minix, The Santa Cruz Operation Inc.'s SCO UNIX, and even at one time Microsoft's Xenix. Variants of UNIX have run on just about every computer architecture, from Intel 8086-based systems to multithousand-node supercomputers from the likes of Cray Research Inc. and Thinking Machines Corp. This incredibly widespread deployment of UNIX alone is a testament to the operating system's success.

The widespread deployment of UNIX also makes it important to carefully define the term "UNIX," particularly when comparing UNIX to other operating systems. The history of UNIX is a sordid one, and a number of standardization bodies have addressed the issue of defining a "standard" UNIX, most notably the IEEE POSIX standards developed to unify AT&T's System V UNIX and the University of California's BSD UNIX (see Figure 2).

Today, the X/Open Group is charged with the branding of UNIX systems based on the Single UNIX Specification.

Although implementations differ, most variants of UNIX are based on a monolithic kernel, with most of the operating system and environment running in kernel mode and providing system calls to user programs. Such an organization is not easily modified or ported, but it is important to realize that the UNIX standard does not specify its organization, only its interface and functionality. It would be possible to have a UNIX branded operating environment on top of a microkernel architecture, such as Softway Systems Inc.'s OpenNT product. OpenNT extends the Windows NT POSIX subsystem to include 1003.2 compliance, ISO/ANSI C language bindings and full XPG4 UNIX95 compliance (see "OpenNT – The Best of Both Worlds?").

UNIX was designed as a multiuser time-sharing system for programmers. That is, it was designed for relatively sophisticated users and, as such, was made to be as efficient and programmer-friendly as possible. As a consequence, it is often quite terse and not very user-friendly. The UNIX naysayers often list this as their number one complaint, but user-friendliness was never a design goal of UNIX. UNIX is also synonymous with client/server computing and Internetworking, having grown up with the TCP/IP protocols and the Internet itself, and is the operating system that launched the World Wide Web.

Scalability

Scalability is a term generally used to denote an operating system's ability to perform well over a range of hardware configurations, particularly clustered systems and multiprocessor (MP) systems. Multiprocessor computers are systems that include more than one processor connected via a communication network and which share a single memory area. The advantage to this is that several applications may run concurrently, each on a different processor, instead of sharing a single processor. In addition, a single application that is multithreaded—meaning that it has several processing tasks (or threads) that can execute at the same time—may be distributed across the processors for increased performance. HTTP (Web) servers are an exemplary case of an application appropriate for multithreading, performing transaction-based processing against a set of shared data, and most existing servers are multithreaded. At this stage, most

Figure 2. POSIX Standards

1003.1	System Calls
1003.2	Shell/Utilities
1003.3	Testing/Conformance
1003.4	Real Time
1003.5	Ada Bindings
1003.6	Security
1003.7	Systems Administration
1003.8	File Access
1003.9	FORTRAN Bindings
1003.10	Supercomputing

Web servers are essentially serving data out of databases, and databases are another perfect candidate for multithreaded/multiprocessor execution, also being based on a transaction model with shared data.

Even though UNIX was not originally designed for multiprocessor computing, it has seen years of research and development to function well on MP systems. Today, there are a range of MP UNIX systems available, such as Sun's Enterprise systems, which link up to 64 UltraSPARC processors via the Ultra Port Architecture crossbar switch network, IBM's RS/6000 SP systems with up to 512 PowerPC or POWER2 Super Chip processors via the SP Switch, or SGI's Origin S2MP systems with up to 128 R10000 processors connected via the CrayLink Interconnect mesh network. Even if your installation does not require 128 or 512 processors, the fact that UNIX scales to such systems should make you confident in running it on two-, four- or eight-processor systems.

Windows NT, on the other hand, was designed with multiprocessor computing in mind from the beginning according to the NT design team and the operating system's historical roots. Out of the box, NT supports up to four processors on the Intel architecture (Enterprise edition supports up to eight processors). OEM NT versions have been successfully tested on configurations with up to 28

processors, including DEC AlphaServer systems with up to eight processors. According to NT Product Manager Gary Sharer, NT is limited to a total of 32 processors on all architectures. What is unclear, however, is whether this is a technical or political limit and how aggressively Microsoft wishes to pursue and support large-scale multiprocessors.

One metric that is commonly used to gauge MP scalability is the TPC benchmarks from the Transaction Processing Performance Council. These benchmarks measure transaction processing and database performance, and by comparing results across systems with varying numbers of processors (all else being equal) an indication of performance gain per processor can be obtained. These benchmarks, viewed in this way, are measuring the complete system: architecture, operating system and application. In the best situation, a system with two processors will have twice the performance of a system with only one. However, there are hardware and software overheads that prevent the achievement of linear speedup (see Figure 3).

Clustering is a technique used to provide fault tolerance and high availability to mission-critical systems. The idea is to connect a number of (functionally) identical systems together in a cluster. Software monitors the running systems and should any hardware or software failure occur, the affected systems are "failed over" to a running system. For example, you might have two identical Web servers, one running and the other to be used as backup. The clustering software monitors the running system, and if a failure occurs, the backup will take over. The failover process may involve redirecting HTTP requests, rebinding IP addresses, notifying an operator of the failure and so on.

Most of the major UNIX vendors have provided clustering products for UNIX platforms for some time, including IBM's HACMP for AIX, DEC's TruCluster products for Digital UNIX, SGI's Challenge cluster systems, Sun's Enterprise Cluster systems based on the Scalable Cluster Interface, and Thinking Machine's GlobalWorks technology. Last year, Sun announced the Ultra HPC servers, which combine SMP server hardware with Platform Computing Corp.'s Load Sharing Facility for distributed, parallel and clustered system support. LSF's Multi-Cluster component even enables multiple clusters to share workloads. Specifically for clustered Web sites, Resonate's Dispatch allows a cluster of Internet servers to act in concert as a single system. Dispatch provides high availability by routing requests around failed servers and, according to Resonate, achieves near-linear performance improvements.

There are clustering solutions for Windows NT as well, one being developed by Microsoft and several others from OEMs. Code-named "Wolfpack," Microsoft's Cluster Server for NT was released last year. Tandem Computers, maker of MP and cluster systems for both UNIX and NT, has developed an NT clustering solution as well, based on its ServerNet technology, a packet-switched crossbar interconnect system area network architecture. The Tandem S-series NT Server clusters make use of ServerNet technology via Peripheral Component Interconnect (PCI) cards and the Cluster Availability Solution (CAS), which is fully compatible with Microsoft Cluster Server APIs, Tandem says.

OpenNT – The Best of Both Worlds?

Many vendors have attempted to bring Microsoft Corp. Windows to the UNIX desktop, with products such as Sun Microsystems Inc.'s WABI and numerous multiuser NT products based on Citrix Systems Inc.'s WinFrame product. Softway Systems Inc., on the other hand, is doing just the opposite by developing a UNIX branded operating environment on top of NT. This is the only product that provides both NT and UNIX on the same desktop, without the performance penalties of emulation or distributed access.

Such a product promises the development environment and tools (including a Motif development environment) of UNIX with the commercial application availability and standardization of NT. OpenNT is POSIX 1003.1 and 1003.2 compliant, and Softway Systems is aiming for UNIX 95 branding. However, according to the company, such branding is low on customers' wish lists.

Included are many of the components common to UNIX systems, including common shells, an X11R6 X server, multiuser support (with a telnet server), shared memory with named pipes, color curses and a complete development environment. OpenNT essentially provides a UNIX front end to the Win32 subsystem—the compiler back-end is Microsoft's Visual C, user security and management is based on the NT security object, and BSD sockets are mapped to the Win32 Winsock, for example.

A number of popular UNIX tools have been ported to OpenNT as well, such as Perl, GNU's `gcc`, `tcsh` and the Apache Web server.—*iw*

Figure 3. TPC-C Results

System	Processors	OS	Database	tpmC	\$/tpmC
Sun Ultra Enterprise 6000	8 (250 MHz)	Solaris 2.6 UltraSPARC	DB2 Universal	15,694.34	90
Sun Ultra Enterprise 6000	16 (250 MHz)	Solaris 2.6 UltraSPARC	Oracle 7	23,143.65	119
Sun Ultra Enterprise 6000	24 (250 MHz)	Solaris 2.6 UltraSPARC	Oracle 8	31,147.04	109
Compaq ProLiant 5000	2 (200 MHz)	NT 4.0 Pentium Pro	Oracle 7	5,008.47	87
Compaq ProLiant 5000	4 (200 MHz)	NT 4.0 Pentium Pro	Oracle 7	6,393.20	109
Compaq ProLiant 5000	4 (200 MHz)	SCO UnixWare 2.1 Pentium Pro	Sybase SQL 11	8,311.43	96

The tpmC results measure total transaction throughput, and \$/tpmC indicates the price/performance ratio. While absolute comparisons cannot be made due to the different databases used, these results indicate that the Enterprise systems scale from eight to 16 processors with a speedup factor of 0.74 and to 24 processors with a factor of 0.66, while the ProLiant systems scale from two to four processors with a speedup factor of 0.64. The throughput for the same ProLiant system with UnixWare and Sybase SQL is 30% higher than with NT and Microsoft SQL, while the price/performance ratio actually decreases.—iw

Security

One primary area of responsibility for any operating system is security. There are various forms of security an operating system can provide, including process security, which provides protection for processes running in the system; user and data security, which provides protection for user-owned data and objects in the system; and system security, which provides protection from sources external to the system, such as other systems on a network.

Both UNIX and NT make use of a virtual memory model to isolate processes' memory address space, which allows for operating system and hardware support of process security. In such a model, memory access violations can be detected and prevented. Beyond that, the process protection mechanisms in the two operating systems differ. NT, with its object design, views processes, threads and other resources as objects, and the entire security model of NT is based on the notion of object security. Because many parts of NT are based on a client/server model, with objects moving from place to place, it is more efficient and reliable to attach security and permissions information (Access Control Lists) to the objects themselves, rather than rely on a single centralized security scheme. UNIX, on the other hand, relies for the most part on the kernel for security enforcement among system components, and almost all internal interaction is through the kernel.

For user and data security, the most often quoted specification is from the U.S. Department of Defense National Computer Security Center "Orange Book," or Trusted Computer System Evaluation Criteria. This sets forth criteria against which computer systems are evaluated and classified. The most common classification is C2, which is a prerequisite for products sold into most U.S. government installations.

The criteria required for C2 classification include:

- **Identification** – All users of the system must identify themselves to the system.

- **Authentication** – All users of the system must be authenticated (password sufficient).

- **Access Control** – All resources must be owned by a user, and users must have discretionary control over who can access that resource and how it is accessed.

- **Auditing** – All security-related events and user actions must be auditable by authorized users.

Both UNIX and Windows NT conform to the C2 criteria for security. However, whereas UNIX has been around much longer to be pounded on by hackers, NT is relatively untried. The number of new vulnerabilities found in NT far outpaces those found in UNIX, simply because UNIX has been under attack much longer. This does not mean that UNIX is rock-solid though; the last report I saw listed five new security bugs for NT, and two new bugs for Solaris, Sun's UNIX implementation. There will continue to be security holes found in both UNIX and NT, which will continue to be patched by their vendors.

Networking

Windows NT is the first personal computer operating system from Microsoft that builds networking support into the operating system, rather than layering it on top. Out of the box, Windows NT supports peer-to-peer networking services such as file copy, electronic mail and remote printing, based on Microsoft's server message block (SMB) protocol and NetBIOS API. In addition, other transports, such as TCP/IP, IPX/SPX, DECnet and AppleTalk may be loaded and unloaded from the NT system. Windows NT supplies APIs for routine I/O, browsing, named pipes, mailslots, sockets and remote procedure call (RPC).

Most versions of UNIX, on the other hand, use TCP/IP as the default transport and have since its introduction with 4BSD more than 20 years ago. Berkeley UNIX also pioneered the socket network interface design, which provides a simple and reliable framework for network communication. Like NT, a


number of other transports may be loaded into the UNIX kernel. In addition, most TCP/IP networking support protocols and services common on the Internet and Web come from the UNIX world, such as the Network File System (NFS), the Domain Name System (DNS) and Berkeley Internet Name Domain (BIND), the Simple Mail Transport Protocol (SMTP) and Network News Transport Protocol (NNTP). The main observation here is that UNIX and UNIX developers have been involved with client/server computing for a long time, whereas client/server computing is relatively new to the PC LAN. If NT is to embrace and support enterprise-scale client/server computing, as seems the trend, then Microsoft has difficult decisions to make. Full client/server computing, with serious security support, will require intricate administration, and this is one edge (ease of administration and use) that Microsoft does not want to lose to UNIX vendors. At the same time, the trend for UNIX vendors has been to make administration and use more simple and more user-friendly. For example, Version 2.6 of Sun's Solaris allows most installation and administration functions to be performed through a Java-enabled browser (in addition to traditional administration techniques).

Applications

By now it should be clear that, as far as Web server platforms go, both UNIX and NT have all the basic functionality required. One area where the differences may be most pronounced is applications. While there is a large, and growing,

number of vendors supporting Web applications on both platforms, there is also a large number of applications that are only supported on one or the other. In addition, because it is common to feed Web servers with existing data from existing applications, some consideration must be given to non-Web-specific applications as well.

If you want to provide a Web interface to your product catalog, for example, and that catalog is stored in a UNIX database, then you will probably have an easier time connecting a UNIX Web server to this data than an NT Web server if only in terms of maintenance and support. On the other hand, if you want to put Microsoft PowerPoint presentations on the Web, then an NT server will make the job much easier.



Then there is the development environment. While both NT and UNIX provide robust development environments, many developers have strong convictions about their favorite and may only be experienced in that environment. Moreover, even though the NT platform supports Sun's Java, and there are efforts to support Microsoft's ActiveX on UNIX, the choice of platform may be an implicit embracing of one computing vision over the other. In the Microsoft camp sits the NetPC with the Zero Administration initiative, in which essentially a PC is made physically secure and is required to support a number of LAN protocols for configuration and administration. In the UNIX camp (spearheaded by Sun, Oracle Corp. and Netscape Communications Corp.) sits the NC, or Network Computer, in which essentially a sophisticated data terminal is required to support a number of Internet protocols for configuration, administration and universal data access. These visions are not mutually exclusive, but they are not equivalent either.

The bottom line is there is little to distinguish the two operating systems based solely on Web server application support. This choice is much more dependent on the environment in which the Web server is to run: the back-end data to be served, the technical support and programming staff, the client environment and the existing infrastructure. Even some of these may eventually become nonissues, based on the industry trend of better integrating these two operating systems with products such as OpenNT.

For Your Information

Helen Custer has written an excellent text on Windows NT implementation details in *Inside Windows NT*, ISBN 1-55615-481-X (Microsoft Press). Also, Andrew S. Tanenbaum's seminal texts, *Operating Systems Design and Implementation*, ISBN 0-13-637406-9 (Prentice Hall), and *Modern Operating Systems*, ISBN 0-13-588187-0 (Prentice Hall), provide accessible explanations of operating system design issues. ➔

BONUS TO SUNEXPERT ADVERTISERS OF THE FOLLOWING ISSUES

JULY ISSUE
SIGGRAPH '98
 Orlando, FL

AUGUST ISSUE
INTERNET EXPO
 Boston, MA

SEPTEMBER ISSUE
COMDEX ENTERPRISE WEST
 San Francisco, CA

IT FORUM
 New York, NY

INTERNET COMMERCE EXPO
 Los Angeles, CA

OCTOBER ISSUE
FALL INTERNET WORLD
 New York, NY

Communicating without a Net

The promise of wireless data communications is simple: free bandwidth. Just buy a pair of wireless Ethernet bridges, put antennas on a couple of roofs, and you can have yourself the equivalent of a T1 data circuit (or faster) between two buildings without spending another dime.



The economics of wireless are so compelling that it's scary. Two radios, antennas and installation cost less than \$8,000. A T1, on the other hand, costs between \$750 and \$1,000 per month in most cities—plus \$1,000 to \$5,000 for the routers you need at each end. That means that the wireless circuit pays for itself in less than a year—usually in less than six months. After that, the bandwidth is just gravy.

And the economics just get better and better as you connect more buildings. Most wireless systems support a hub-and-spoke arrangement, which means you can add buildings by buying a single radio and antenna. Even better, many wireless networks support low-cost PCMCIA cards that plug into laptops, letting users cut the cord and carry their laptops around the office—or outside—without losing connectivity.

The promise of wireless is so seductive that more and more businesses and universities are turning to wireless networking, which is in turn creating a booming business for wireless vendors. There is now even an IEEE standard for wireless Ethernet bridges. But, as my company learned over a very difficult six months, there can be a big difference between the promise of wireless and the reality.

Unlicensed Spectrum

High-speed wireless communications using microwaves has been around since the 1960s. (After all, the initials in that telecommunications giant's name originally stood for Microwave Communications Inc.) But back then, wireless

communications were expensive. Beyond the cost of the transmission towers and radios, businesses that wanted to communicate using microwaves needed something else: a license from the Federal Communications Commission (FCC) giving them the right to transmit at a certain frequency, with a certain power and over a certain geographical area.

For some businesses, radio licenses have always been essential to doing business. Taxi companies and all Federal Aviation Administration (FAA)-licensed pilots need radio licenses to do their jobs. But, for many other companies, the whole process of getting a radio license has been shrouded in mystery. Even if it were simple, few people understood how to do it. As a result, many avoided the technology.

Conventional radio systems have another problem: interference. If somebody is transmitting on your frequency, their signal can easily overwhelm yours (or vice versa). But in the 1970s, a new technique was discovered for transmitting radio waves called frequency hopping, or spread spectrum. The systems minimize interference by transmitting a signal successively on dozens (or even hundreds) of different frequencies. The receiver knows the pattern that the transmitter is using and follows along. Because interference usually doesn't hop between frequencies (and if it does, it usually doesn't follow the same hopping pattern that a particular spread-spectrum user has chosen), whatever interference happens is usually temporary.

As things turned out, the U.S. military knew about spread spectrum for many years: Besides minimizing interference, spread-spectrum signals are exceedingly difficult to trace and eavesdrop. But once

the secret was out, another problem reared its head: Using the technique inside the United States was illegal under FCC regulations because radio applications were licensed to a specific frequency.

Things changed in 1985, when the FCC set aside a portion of the radio spectrum for spread-spectrum systems. Eager to avoid the fiasco of CB radio, the regulations (Section 15.247 of the FCC rules) required that unlicensed radios have a maximum peak output power of one watt and a maximum antenna gain of 6 dBi.

Today there are three radio bands set aside for unlicensed use: 902 - 928 MHz, 2,400 - 2,483.5 MHz and 5,725 - 5,850 MHz. To use the spectrum, all you need to do is purchase a radio that has been approved. In fact, you might already be using the same license spectrum if you have a frequency-hopping cordless telephone such as those made by VTech Communications Ltd. and AT&T.

Networking the Island

I run a small Internet service provider (ISP) on Martha's Vineyard called Vineyard.NET that offers dial-up service to roughly 1,000 users. We first started thinking seriously about radio back in February 1997, when we started building a second machine room. The second machine room was our insurance policy: If we had a fire at the first location, we could continue operations at the second. Our design called for two computers to be connected by a high-speed network connection, with constant disk-to-disk backups between the two. Naturally, we wanted to do it as inexpensively as possible.

The obvious solution was to set up a T1 between machine rooms, but the phone company wanted roughly \$1,000 per month for the connection. Another option we considered was hanging our own fiber overhead. We did a little research and discovered that the cost would have been thousands of dollars for the cable and the engineering plan required by the power company, an undetermined expense for the truck and police detail to put the cable up and a pole rental of roughly \$5 per pole per month. If we ever had a downed line, we would have to start all over. Finally, we considered using a pair of Asymmetric Digital Subscriber Loop (ADSL) modems. Although ADSL manages to drive something like 6 MB/s over an ordinary pair of copper wires, the distance between each of our buildings and the telephone company's home office was further than the ADSL modems could handle. To make matters worse, the modems had a list price of \$4,000 per pair—and you still needed to buy a router for each end if you wanted to connect two networks.

In the midst of our deliberations, we discovered wireless. We were immediately beguiled. The speeds promised by wireless were faster than T1. The price for the first connection was roughly equivalent to the cost of the ADSL modems or a new Cisco Systems Inc. router at each location. But unlike ADSL, we could add locations by purchasing a single radio, rather than a pair.

It was this expandability that really caught our interest. If the wireless system worked, we could use it for more than simply connecting our two machine rooms together. We could use

the wireless network to provide high-speed access to our customers. Small ISPs like Vineyard.NET are facing increasing competition from big companies such as AT&T and Sprint Communications Co., who offer nation-wide dial-up at a cost that is too low to make a profit. Wireless, we learned, has a range of roughly three to five miles, and Martha's Vineyard is pretty flat. This meant that we might be able to use our machine-room connection to offer low-cost, high-speed Internet service throughout the island. If we could build a wireless network on Martha's Vineyard, we might just be able to find a lucrative niche that would let us survive the upcoming ISP shakeout. At least, that's what we thought.

But which wireless system to use? We had no clue. The fellow who first suggested that I consider wireless said he was using equipment manufactured by an Israeli company called Breeze Wireless Communications Inc. We clicked into the BreezeCom Web site. The equipment seemed reasonable and, reassured by the company's 30-day money-back guarantee, we ordered ourselves a \$4,000 starter kit.

BreezeCom

Most of the radios that BreezeCom makes are about the size of a paperback book. Each radio has two antenna jacks, connectors for RS-232 and power and either one or four 10BaseT connectors. The units have six LEDs that display the condition of the wireless connection, the condition of Ethernet and power.


At the heart of the BreezeCom system is a special radio called the Access Point (AP-10). This unit is the hub of your wireless network. We actually ordered the AP-10DPro with a single 10BaseT connector that you patch into your Ethernet hub and two antenna connections. It operates in the 2.4-GHz band.

(A note about the BreezeCom product nomenclature. The AP-10 is available in a variety of different styles. The "D" means that the product is sold without built-in antennas, which are suitable for use in an office or factory environment, but not outdoors. The "Pro" means that the unit is part of BreezeCom's Pro series, compatible with the agreed upon 802.11 standard, which allows for limited interoperability between wireless vendors.)

To communicate with the Access Point you need either a Station Adapter or a Wireless Bridge unit. BreezeCom manufactures three kinds of station adapters. The SA-10 series connects a single Ethernet device to your Access Point. The SA-40 connects four Ethernet devices; it has four 10BaseT connectors and a built-in four-port 10BaseT hub. BreezeCom also manufactures a PCMCIA adapter called the SA-PC that plugs into a laptop and connects your laptop to the wireless network. And finally, BreezeCom makes the Wireless Bridge that will bridge 256 Ethernet MAC addresses across the wireless network to the Access Point.

With the exception of the SA-PC, all of the BreezeCom units can operate with one or two antennas. Although you only need one antenna, it's better to use two, because that way you can use the system's "antenna diversity" feature. Antenna diversity is supposed to solve the problem of momentary physical obstructions in your radio path by creating more paths for the radio signals to travel. With diversity enabled,

the radios transmit all data using both antennas. When they receive information, a tiny circuit inside the radio automatically switches between the two antennas, choosing the antenna that is receiving the stronger signal. Essentially, diversity increases the number of radio paths through the air from one to four (two antennas, each one capable of reaching two receivers).



We ended up ordering an AP-10DPro, a WB-10DPro, an SA-40DPro and an SA-PC. We ordered a pair of omnidirectional antennas for the hub where the Access Point would be located and directional antennas for the WB-10 and the SA-40. BreezeCom supplied us with the radios, the antennas and special low-loss cable for the interconnection. Unfortunately, the cables that BreezeCom gave us were only 20 feet long. This short length created significant problems in placing the radios. The short cabling is a direct result of the limitation on transmission power by the FCC. BreezeCom told us that it was working on special amplifiers that would allow us to use up to 200 feet of antenna cable, but those amplifiers were not available when we needed them. (Using the amplifiers also destroys antenna diversity because the amplifiers force you to use one antenna for transmission and the second antenna for reception.)

We wanted to mount the antennas on our roof, but the short cable forced us to mount the radios high as well. In our first location (my house), we were able to find a place for the radio in a second-floor closet. But in our second location, a 32-foot-high building with 16-foot-high ceilings, we were forced to build a special shelf that we hung from a ceiling to hold the BreezeCom radio, its power supply and a UPS within 20 feet of the antenna. (We couldn't just put the BreezeCom unit in the building's attic, because the attic routinely reached more than 110 degrees Fahrenheit and the BreezeCom units were only rated for an ambient temperature of 90 F. We couldn't mount them outside because their boxes were not waterproof.)

Installing the antennas on our roof was another difficult task. We had to mount a tripod on each roof, drill holes for the cables and make sure that everything was properly grounded. Grounding is key, as the antenna wire basically provides an electrical pathway from your roof to your machine room—just

the thing to guide a lightning strike from an angry storm cloud. Fortunately, we were able to find an excellent reference book on grounding, *The "Grounds" for Lightning and EMP Protection (GLEP)*, Second Edition, by Roger Block, president and founder of PolyPhaser Corp, Minden, NJ. In addition, we bought a pair of gas-discharge lightning protectors, also from PolyPhaser, for added protection. Each antenna took roughly a day to install.

One of the big drawbacks in the way we mounted the BreezeCom radios was that we did not have easy access to the units. At the time, we minimized this problem, rationalizing that if we needed to have physical access to the radios on a regular basis, then we probably didn't want to be using them in the first place. As it turned out, we were right.

The BreezeCom Setup

To set up the BreezeCom units, you attach a serial cable to the unit's proprietary serial port, plug the other end into a laptop and program the unit with a standard terminal emulator. Stepping through the BreezeCom menus, you can set each unit's IP address, its netmask, a few radio parameters and the amount of logging you desire. Security is provided through a password: Set the password, and the unit's parameters cannot be changed without providing the password, although they can still be viewed.

Sticking largely with the factory defaults, we gave each unit an address on our local network, then rebooted the radios. A few minutes later, we were pleasantly surprised to discover that the units were performing exactly as advertised. Sitting at a computer connected to the WB-10, it was easy to access information stored on the network connected to the AP-10. For all practical purposes, the wireless units were bridging the two physical Ethernets, making them one. Wireless was easy!


But shortly after we got things up and running, we started to have concerns about the quality and design of the software running inside the BreezeCom radios.

At Vineyard.NET, we monitor the throughput of our network with the Simple Network Management Protocol (SNMP) using a piece of free software called MRTG. SNMP doesn't have much security. The only way to prevent other people from monitoring your network is to change your SNMP community from "public" (the factory default on many systems) to some other value. Essentially, your SNMP community is a password. BreezeCom supports SNMP, but the BreezeCom units didn't allow us to change the SNMP community. When we complained to the company, BreezeCom recommended that we disable TCP/IP to protect the units from a network-based attack. We didn't want to lose the ability to monitor the radios, so we compromised by programming our routers to prevent SNMP connections to the BreezeCom radios from outside our local network. This provided some protection, although it did nothing to protect us from a local attack.

After we reconfigured our SNMP software to use the "public" community, we discovered that we were still out of luck. BreezeCom didn't properly implement the particular SNMP values required by our network-monitoring software.

We later learned that BreezeCom only supports its own custom Management Information Base (MIB).

Another concern that we had was preventing possible theft of service. We didn't want somebody to be able to get free access to Vineyard.NET by simply going out and buying their own BreezeCom radio. Unfortunately, BreezeCom didn't give us much in the way of protection. The sole means for preventing unauthorized access was by setting a configuration variable called the Extended Service Set ID (ESS ID) to be something other than the default value. The ESS ID is designed to allow more than one wireless network to interoperate within range of one another without interfering with one another. Each wireless unit will only speak to other units with the same ESS ID. But the ESS ID is designed to prevent cross-talk, not provide real security. In practice, there were many ways for a potential attacker to learn the ESS ID we were using.



We had been using the BreezeCom units for about a month when we installed a used terminal concentrator (a Bay Networks Inc. Annex 4000) on our internal network. A day later, the Access Point stopped working. We had no idea what was wrong. The system had worked perfectly until that point. We assumed that the software had crashed. Not knowing what else to do, and it being a Sunday, we got a ladder and rebooted the Access Point. Things worked well after that for a few days, but then the Access Point crashed again.

A telephone call to BreezeCom revealed that we were not running the most recent version of the Access Point firmware. BreezeCom recommended that we upgrade. But the upgrading process was frightening. To upgrade the BreezeCom unit, we were told to send it two TFTP commands: The first command told the unit to enter programming mode; the second command uploaded the new firmware. What was frightening was the fact that the unit performed no sanity checks on the uploaded binary to make sure that it was correct, and there was

no security to prevent anyone else on the Internet from reprogramming our BreezeCom units without our permission. Furthermore, BreezeCom warned us against using the UNIX TFTP command to do the programming. Unless we used a particular TFTP program for Windows 95 computers, the timing might not work out perfectly, and we might end up rendering our Access Point unusable. And BreezeCom told us to never, never, never attempt to upgrade the firmware over the air.

Unfortunately, we didn't have a Windows 95 computer to reprogram the Access Point. BreezeCom said it would send us a replacement unit but suggested that we try using the UNIX TFTP command, just in case it might work. Sure enough, the timing didn't work out. We sent BreezeCom back the dead unit.

We received a replacement AP-10DPro with the new firmware a few days later, but it had the same problem. Then, our WB-10DPro started malfunctioning as well.

Over the following month, we worked closely with BreezeCom trying to isolate the cause of the failures. BreezeCom had us connect a personal computer to the serial port on the Wireless Bridge to capture all the information that the unit would log. (The system kept some log information in memory, but this information was wiped out every time the machine crashed.) Then, each time we had a failure, we sent the captured serial port output to BreezeCom's engineers. The files contained numerous warning messages complaining about invalid ARP packets. My guess is there was a bug in the BreezeCom TCP/IP stack, possibly being tickled by our Annex 4000.

BreezeCom didn't seem to be prepared to debug its systems in the field. After a few weeks, the company's engineers said that they had many hundreds of units in the field, and we were the only customers reporting problems. But our units were crashing. After many frustrating weeks, we finally gave up.

To be fair to BreezeCom, the company gave us no difficulty when we wished to return roughly \$5,000 of equipment several months after the 30-day money-back guarantee period had expired. Within a few weeks, we had our money back and were ready to try again with another vendor. But this time, we were going to be informed consumers.

The Wireless Directory

Before we went shopping for our second wireless system, we searched the Internet for all the information that we could find. The most useful service that we found was a directory of wireless products maintained by Barry McLarnon, a ham radio operator and a researcher of digital radio broadcast systems at the Canadian government's Communications Research Center (see <http://hydra.carleton.ca/info/wlan.html>). The directory includes a comprehensive list of products that operate at the 915-MHz, 2.4-GHz and 5.8-GHz bands, as well as reviews, articles and links to vendor Web sites.

One of the decisions that we made at Vineyard.NET was to change our primary frequency from 2.4 GHz to 915 MHz. Our reason had to do with the trees. Vendors are excited by the higher frequencies because they offer better noise immunity in crowded urban environments, which means that most customers are less likely to have interference. But on rural

Martha's Vineyard, lower frequencies get better propagation around the rolling countryside. Lower frequencies also have a better chance of shooting through tree canopies, which are all too common here.

Restricting ourselves to 915-MHz vendors, we quickly had a short list of three companies with products that we thought would be acceptable. So we made up a long list of questions, most of them having to do with security, then called each vendor to find out their answers.

Why focus on security? One reason is that I specialize in this area. But security was also the first area in which we discovered problems with the BreezeCom units. Security is very hard to do. Our theory was that if a company got its security correct, it would probably get the other features correct as well.

When we started making phone calls, we discovered something very annoying about the wireless directory: Most of the prices were wrong. McLarnon says that the directory is very big and it is difficult to keep up-to-date. One of the things that makes the job particularly difficult is that some prices include the cost of the antenna and/or cables, while others require that you purchase this equipment separately.

Ultimately, we decided to go with a company called C-Spec Corp., Dayton, OH.

C-Spec

C-Spec's wireless network system is called OverLAN. The radios are built from a low-profile industrial PC chassis equipped with a 486 CPU, an Ethernet card, a Lucent Technologies Inc. WaveLAN card, a proprietary card containing flash RAM and an interface for the front-panel lights. The system can transmit up to 2 Mb/s. C-Spec also makes a PCMCIA card and an ISA card. Any OverLAN unit can operate as either a hub or a spoke.

The C-Spec radios have two modes of operation. Like the BreezeCom units, they can operate as a wireless Ethernet bridge, connecting many Ethernet segments together. Alternatively, the units can operate as true routers, linking multiple TCP/IP, Novell Inc. or AppleTalk networks in an intelligent, manageable fashion.

C-Spec supports multiple layers of security. The SNMP community can be changed. The system also allows access control lists, so that SNMP commands can only be accepted from particular IP addresses. The over-the-air traffic is encrypted with the U.S. government's Data Encryption Standard (DES). And, unlike the BreezeCom units, C-Spec supports remote software upgrade, even over the air.

Getting the C-Spec units up and running was a little trickier than with BreezeCom. Before they can be made operational, C-Spec requires that its units be programmed using SNMP and a special configuration application that the company has created. The application runs only under the DOS operating system (and not under Windows 95). We spent two days looking for an old 386 computer on which to run the program. The configuration program worked as advertised but occasionally crashed. C-Spec promises that it will have a new, Windows 95-based program shortly.

Since they have been installed, the C-Spec units have operated flawlessly. They support the full SNMP MIB II, allowing us to integrate them easily with our current network management software. And the large number of front-panel LEDs (showing the quality of the wireless signal as well as congestion on both the wireless and Ethernet networks) give us a good idea of what is happening with the radio simply by looking at the unit. Our only complaint with the C-Spec units is their noisy fans (two fans in each radio). The company says that it used noisy ball-bearing fans because they are less likely to fail than quieter units. Fortunately, C-Spec supplies cables that are 50 feet long, so we were able to install the units in our basement, out of earshot.

C-Spec claims its systems will operate up to three miles (creating a cell six miles across). The company sells an optional amplifier that boosts the range to 10 miles.

Is Wireless in Your Future?

The outlook for wireless looks good. Systems are getting faster and faster. C-Spec, for example, claims a new version of its OverLAN system offers 10 Mb/s. And a growing number of vendors support the 5.8-GHz band, which is good news for people in crowded urban areas.

The IEEE has passed the 802.11 standard, which theoretically allows interoperability between different spread-spectrum wireless products. BreezeCom has been very active in the standardization activities. You can find out more about the standard at <http://www.breezecom.com/80211.html>.

Nevertheless, it's not clear just how much difference these standards will actually make. Many vendors say that 802.11 will not be as important as many users anticipate. The standard does not guarantee compatibility or even interoperability between different products. Furthermore, many wireless systems may produce higher throughput or have more management features when using the company's own proprietary protocols, instead of the standard.

If you need to connect two sites that are within three miles of one another, you have access to your roof and you're not surrounded by taller buildings, wireless might be a very cost-effective way to build your network. Alternatively, if your office is located on a high spot, and there are engineers who want to work from home who can see the office building, wireless might be an easy way to give them free high-speed access.

As the cost of wireless PCMCIA cards drops (expect \$500 cards this year), many universities and offices are likely to use indoor wireless systems as a flexible alternative to pulling Ethernet wire. Combined with a Dynamic Host Configuration Protocol (DHCP) server, these wireless systems can provide instant access and transparent roaming for laptop users, even when moving between buildings.

If you do decide to go wireless, it's important that you have some kind of backup between the two facilities. After all, the real danger of the unlicensed spectrum is that it is not guaranteed. If some interference comes along, it can shut your system down without any legal recourse on your part—although you can always move to another frequency. ➡

Clustering

No matter how big a server is, no matter how fast the processor or how much memory is installed, there are still numerous bottlenecks that affect performance and represent single points of failure. Clustering techniques alleviate these problems for high availability and load balancing of service.



Clustering is a technique whereby several functionally equivalent objects work in concert to perform some task that could be performed by any one of the objects alone. The idea is to enhance task performance by using multiple objects. The enhancement may be quality of service, speed of service or reliability of service, and there are many types of objects that may be clustered in such a fashion. For example, RAID subsystems are a type of clustering for storage devices that are used to speed storage access, increase reliability and increase availability of stored data. Multiprocessors are a form of clustering for processors that can be used to increase performance.

Clustering, however, is usually used to denote clusters of servers providing some service, such as database access. A cluster of servers can increase performance by distributing service requests across the entire cluster (instead of a single server servicing all requests), and can increase reliability by eliminating single points of failure (if a server fails, the other servers continue to service requests). There may still be single points of failure in the service, however. For example, clustered database systems usually run a database engine on each of the clustered servers, but the database itself is stored on a single shared storage device to which all the servers have access. In this case, the shared storage device still represents a single point of failure.

For a Web site, clustering can offer the same advantages of increased performance and high availability. Even if you are running your site on the latest and greatest multiprocessor with multithreaded server software and gigabytes of memory, you face limitations on the scalability of service in the form of a single disk controller or network interface. It doesn't matter if the system can process 10,000 requests per

second if the network card can only handle 5,000. In addition, clustering can provide high availability of the Web site; if one or more servers fail, requests can be routed to other running servers. Imagine if your Web server went down at 3 a.m. on a Saturday. How long would it be before you knew it, and how long until you could repair it? There are also cases when a system must be brought down for maintenance; with a cluster, maintenance can be performed without disrupting service.

Virtual Addresses

When constructing a cluster of servers for Web services (or other TCP/IP services) one hurdle is addressing. Every node in a TCP/IP network has a unique IP address. Name and directory services map address names to IP addresses. So, for example, when you enter a URL into your Web browser, the node name part is resolved to an IP address, the IP address of that server, and this IP address is used to communicate with the server. Because each server in a cluster is a separate network node, each has a unique IP address. Whether you wish to distribute Web site service over a cluster or use that cluster to provide high availability, addressing is a problem, because you want to provide a single address name for the site. There are several ways to overcome this problem. One approach is to use more intelligent naming services. When a user enters a URL, the browser looks up the IP address of the Web server using some naming service, usually Domain Name System (DNS). There is an entry in the name service that binds the host name to the IP address of the Web server. Because it is this binding that browsers use to find your Web server, and this binding associates a single node with the name, we

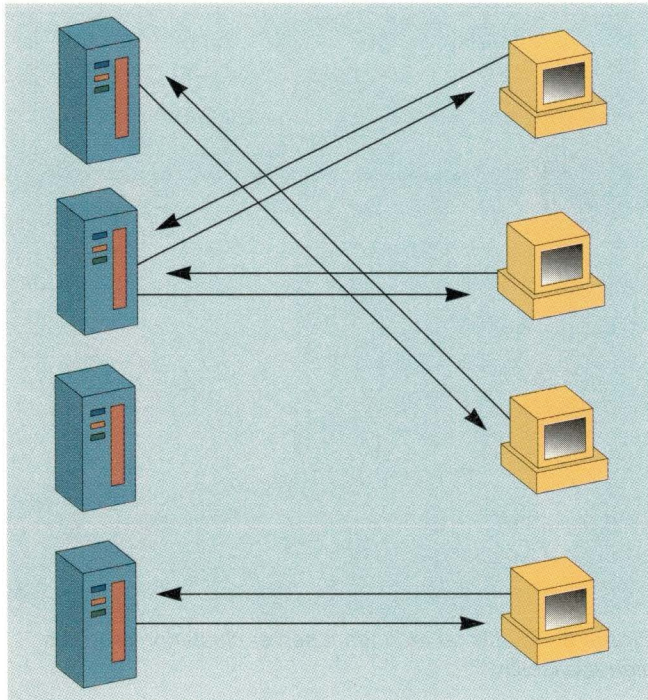


Figure 1. Virtual addressing via hybrid naming services provides coarse-grained load balancing.

cannot directly set up a cluster for Web services. What we want is the name service to associate a name with a set of IP addresses, namely the addresses of all the servers in the cluster, and there are hybrid DNS servers available with this functionality.

When a name service lookup is performed on a name, the name service must return a single IP address for that name. Thus, the name service must decide which of the set of IP addresses to return. Because name services are usually distinct from other types of services such as Web sites, they really can't make an informed decision about which address to return. They can return addresses in round-robin fashion, however: a different address each time a lookup is performed. This has the effect of spreading load across all the servers in the cluster, but this load distribution is at a coarse level (see Figure 1).

Because many name services, and DNS in particular, are distributed services, they make use of caching to increase performance. When a browser looks up a name, it usually contacts a local name server. This server contacts a higher server in the hierarchy and so on, until the server responsible for that name is contacted. The IP address returned is cached by the browser as well as the name servers so that the next time the name is looked up, the address is close at hand. This improves name lookups, but it thwarts our round-robin addressing scheme. Only the lookups that actually reach the domain name server will be given different cluster addresses, while cached lookups will use the same address.

You can exercise some control over the caching of name bindings by setting the binding's time-to-live (TTL) very short, which is the time that other name servers are allowed to cache bindings, but this still doesn't provide a very fine grain of control.

Virtual Servers

Another problem with the name server approach is that it is not very intelligent about which IP addresses to return. It assumes all servers in the cluster are identical, and it doesn't know anything about the status of those servers. You may want to return the IP address of the server with the lowest load at the moment (and at least not return the address of a failed node), or you may want to return the IP address of a server that is geographically closer to the client making the request. You may want to divide your Web site functionally among the cluster servers, with, say, static pages on one set, dynamic pages on another and a search engine on a third set. You want to choose cluster servers based on server load and availability and client request and client location, which is difficult to do with hybrid name services.

Another approach is to bind the Web site name to a single node, but a node that can make intelligent decisions about things like load and location, and redirecting client requests to the appropriate server. This node essentially acts as an application intelligent IP gateway, using TCP connection hops to "redirect" incoming HTTP requests to the chosen server. When a user inputs a URL, the browser looks up the IP address of the server and receives the IP address of the gateway. The browser then sends a request to that IP address. When the gateway receives the request it may look at what type of request it is, where it came from, or the current status of the servers in the cluster. The gateway makes a decision as to which server should handle the request, then performs a connection hop to that server. The connection hop preserves the original source of the request, so the server sees the request as coming from the client and not the gateway, (see Figure 2).

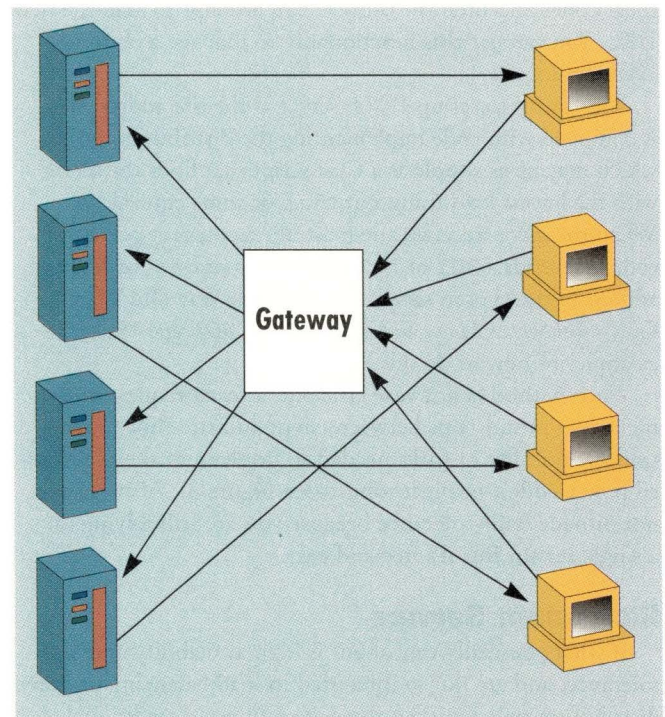


Figure 2. Virtual servers provide fine-grained load balancing and can implement other distribution policies as well.

There are a number of such products commercially available, including F5 Labs Inc.'s BIG/ip product and Resonate Inc.'s Dispatch product. The former, billed as a "network controller," is a hardware product that continuously monitors server and service availability, load balancing TCP/IP (including HTTP) requests across the cluster. To provide fault tolerance and high availability of the network controller, F5 Labs offers BIG/ip2. This is a configuration in which two BIG/ip controllers are linked; one a primary and the second a backup. The backup unit monitors the primary and takes over in the event of a failure.

Resonate's Dispatch, on the other hand, is a software product in which one or more servers are designated as "scheduler" nodes that implement distribution policies and perform the connection hop. Dispatch scheduler nodes first parse incoming requests to perform resource-based scheduling on the service, host, port and path name of the URL. This scheduling produces a set of servers that are eligible to service the request. Then, the scheduler applies load-balancing rules to choose the server to hop to. Load balancing may be performed round-robin, or by criteria such as CPU load, number of open connections or network latency.

HTTP Redirect

You can avoid modified naming services or manipulating TCP by making use of redirection functionality built into HTTP. In order to handle the possibility of resources changing location, HTTP provides several status codes to indicate the relocation of a resource. These include the 301 Moved Permanently and 302 Moved Temporarily response codes, which both include a URL for the new location. When a browser makes a request, the server can respond with one of these codes and then the browser can attempt to load the new URL. You can use this functionality to manage a cluster of Web servers.

The idea is to set up a Web server at the site address that is equipped with code implementing the distribution policy, which may be as simple as a CGI script that finds the server with the lowest load using `ruptime` or some equivalent. When the server receives a request, it executes the policy code and sends a 301 or 302 code to the browser with the address of the chosen server. (See "Which Way Did They Go?," *WebServer Magazine*, May/June 1996, Page 19, for examples of how to do this.)

This method is not very efficient, as each request must make two round-trips between client and site (not to mention the number of forks needed to implement the mentioned policy code), nor is it robust (see Figure 3). Also, it does not provide fault-tolerance because you are still relying on a single server. But it's free and easy.

Redundant Servers

If what you really care about is high availability and fault tolerance, and are not so interested in load balancing or other distribution policies, then there are numerous traditional products for server fault tolerance commercially available. By providing redundant components and by monitoring the




Figure 3. Traditional redundant server clustering provides high availability.

running system, backup components may be brought online in the event of failure. For servers, this can take the form of redundant components in the server chassis, or redundant complete server systems. An example of the former is Integrix Inc.'s iNS200, which can include a redundant motherboard (including up to four Ethernet interfaces per board) and either RAID or JBOD for redundant storage. A diagnostic module monitors the system boards, power supplies, fans and other components. Most major vendors have solutions for the latter, including Microsoft Corp.'s Cluster Server, IBM Corp.'s HACMP and Sun Microsystems Inc.'s Enterprise Cluster offerings. These products typically connect the redundant servers to the primary via SCSI or serial connections, and software monitors the running system to detect failure. In the event of failure, the redundant system "wakes up" and takes over. Because the redundant systems are functionally equivalent, including IP address and running software, the failover is seamless (or should be).

Intelligent Clients

Another possibility if you are developing the client to access the cluster, is to do what Netscape Communications Corp. did. When you click on the icon in the Navigator browser, while it looks like you're going to `http://www.netscape.com`, the client is actually accessing a randomly chosen server from its cluster. This acts very much like the hybrid DNS scheme mentioned above, but, in this case, the client is choosing the IP address to use rather than the name service. In fact, this scheme has other advantages as well. Because the client knows what request it is making and where it is located, it may be better suited to decide which cluster server to send the request to. However, this requires specific knowledge of the cluster and service, and so is most appropriate to intranet services. Also, this type of solution also does not automatically provide high availability and load balancing. ➡

A Good Web Server Platform at the Right Price

Ian Darwin is a consultant, developer and trainer specializing in UNIX operating systems and the Java programming language.

A key decision in setting up and maintaining a Web site is the choice of operating system and hardware platform. Now, before you say, "I'll just use Windows NT" or "I'll just use Solaris," think about this: The operating system needs to be affordable, reliable, secure, fast, scalable, have Web server software available for it and have a list of satisfied users.



OpenBSD, the newest variant of 4.4BSD or "Berkeley UNIX," meets the criteria, plus, it runs on nine hardware platforms. Of course, every operating system vendor wants you to believe that its OS is best (for all purposes). Let's leave discussion of commercial UNIX systems' and NT's (in)security to the Bugtraq list (<http://www.geek-girl.com/bugtraq/index.html>), and here concentrate on how OpenBSD fits the above criteria.

OpenBSD Installation

The OpenBSD installation has some very nice features. First, on four architectures (Sun Microsystems Inc. SPARCs, Intel Corp. i386s, Digital Equipment Corp. Alphas and Advanced RISC Computing, or ARC, machines), there's the ability to boot directly from one of the two CD-ROMs in the kit. For example, on my SPARCstation IPX, I just said `boot cdrom`. My PC did not support this option, so I had to boot from a floppy. The floppy image is provided and can be copied to floppy either using `dd` or using a DOS program `rawrite`, also provided. When you boot from the floppy, you first see a UNIX-style boot listing. After the kernel is initialized, the version of `init(8)` on the floppy asks which shell you want; hitting return gives the Bourne shell (which is really PDKSH). At this point, you type `install` and you're off to the races.

The `install` script asks the usual questions about which boot disk to install on and lets you run its `fdisk` and `disklabel` programs to set them up. You probably know that on a PC, `fdisk` breaks the disk into larger

chunks; on UNIX, a second program assigns individual partitions within the UNIX `fdisk` partition. At least one is required for swapping, because UNIX gains some efficiency by swapping directly to a disk partition instead of into a file as, say, Windows 95 does. UNIX reserves swapping to a file for when you don't allocate a big enough swap partition.

The one trick to this step is getting the sizes and offsets of the partitions exactly right. It may be easier (and safer, if you are sharing your boot disk with another PC-based OS) to use DOS `fdisk` to create the partition, then change it to an OpenBSD partition using `disklabel`. Please read the accompanying installation documentation for the gory details on disk partitioning.

The `install` script also lets you set up the network, which is necessary only if you are using it to install. Yes, in fact, you can install with just a boot floppy and a network card. The main install files, of which there are five, plus a copy of the generic kernel, can be installed from CD-ROM from a local disk in any of several formats—including DOS/Windows FAT, or from a network via NFS, FTP or even HTTP.

Indeed, if you are patient, you can do your entire installation over the Internet, given a copy of the boot floppy and a direct connection. (Don't try this at home, kiddies, if your Internet link is 56 Kb/s or under. It might work, but you'll probably die of old age first.)

I didn't get to try the network install, but on one machine I did have to copy the install files to a Windows partition and run them from there. It worked like a charm!

Apache – No Muss, No Fuss

First, let's look at software, because without it you won't be able to use OpenBSD as a Web server platform. The most popular Web server on the Internet is the free-software Apache server (<http://www.apache.org>). After you've installed OpenBSD, copied its "ports" tree from the CD to disk and have a connection to the Internet, it's just a matter of typing

```
cd /usr/ports/www/apache
make install
```

and you've got a recent version of Apache compiled and installed, ready to configure. No muss, no fuss.

I didn't test either Netscape Communications Corp. or Microsoft Corp. Web servers because neither is readily available in a form usable on OpenBSD.

There is other software readily available for Linux, and OpenBSD has a good "compatibility" mode for Linux, SVR4 and other systems. This compatibility mode is not system-call mapping (with the overhead that that implies) but a method of loading a different set of shared libraries and system-call entries.

For example, Figure 1 shows a 43-KB GIF image of two Netscape Navigator 4.0.x screens. Both are displayed on a Solaris workstation. One is Netscape running locally (Solaris version), the

other is the Linux Navigator in OpenBSD's Linux emulation. View them side by side and see if you can tell which is which.

What about Perl, the language used by many Webmasters for CGI scripting? No problem! OpenBSD is one of the few systems that ships with Perl preinstalled, ready to use. And what about Java? Of course you don't need Java support just to install applets on pages that you serve; they are just data files to the server. But you might want to develop your own applets.

Or you might want to run "servlets," which are Java fragments that blend into your Web server, avoiding the overhead of starting a new process each time a user posts a form. But Java is a new language: Some love it, some leave it. If you're in the former camp, you'll be happy to know that there's a FreeBSD port of Sun's Java Developer's Kit 1.1, which works reasonably well (that compatibility mode, again). There should soon be an OpenBSD native support. Of course if you're in the latter camp, you won't care beans about Java. In either case, OpenBSD delivers the goods in terms of server software support.

Cost and Support

Windows NT Server costs more than \$500 per server (depending on where you shop and how many client

licenses you buy). Solaris Server Edition also costs more than \$500 per CPU. OpenBSD costs \$30—yes, that's thirty bucks—if you want the full distribution on CD-ROM. If that's too much (and if it is, how did you ever afford a PC to run it on :-)), you can download OpenBSD over the Internet. In either case, there's no per-CPU licensing. So if you add a second or third server, or decide to run it on your notebook PC (having given up on Windows 95 after having to reinstall it for the 95th time), there's no cost.

However, the OpenBSD Project's main source of funding is from sales of the CD-ROM, which pays for development, timely patches from its Web site and the like. So if you're using OpenBSD on dozens, hundreds or thousands of machines, you really should make a contribution to the project—cash, machinery or software (your faithful scribe ported a kernel-level device driver for the Connectix QuickCam from another UNIX platform and donated it to the project. Why? I like [some!] WebCams, I like QuickCam and I had the hardware handy).

As you will have gathered, OpenBSD is free software, and there is no guarantee of support. However, the source is in the hands of many developers. Peter Honeyman of the Center for Information Technology Integration, an applied research and development center at the University of Michigan, chose OpenBSD as the platform for a large research project. He says, "we run OpenBSD because it runs AFS [the Andrew File System, an alternative to NFS], has had a lot of attention paid to security vulnerabilities, is really UNIX and has a friendly and open developer community."

As an example, while my Cyrix Systems Pentium 166 system works fine off the CD-ROM distribution of OpenBSD 2.2, my Dell Computer Corp. notebook's ATAPI CD-ROM wasn't recognized. An email to the mailing list brought a reply from a developer in Europe, including a patch, within about 10 minutes. How many commercial OS vendors even read their email, let alone respond the same hour or the same day?

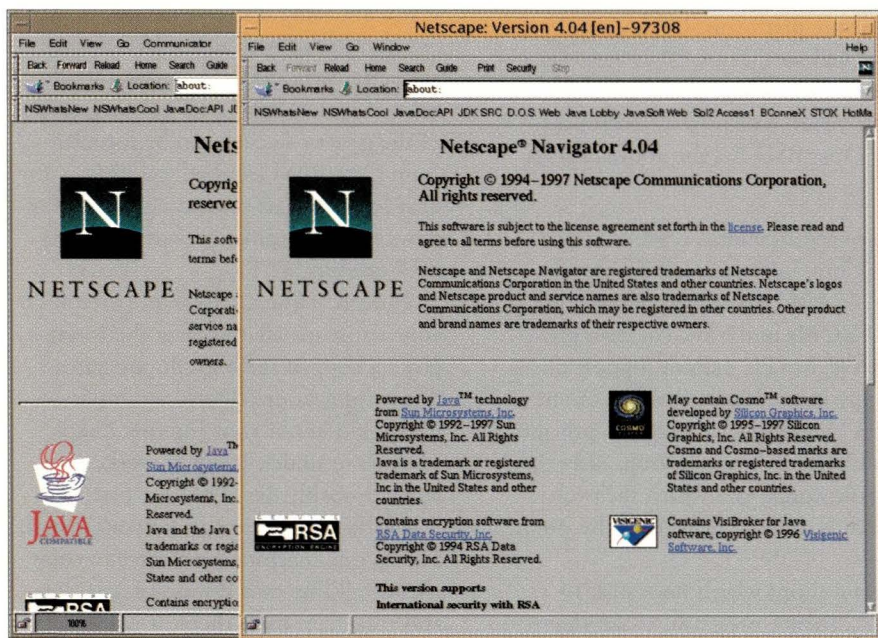


Figure 1. Can you tell the Solaris from the Linux version?

For the Security Wary?

Today's Internet has its full share of vandals and crackers. To stay alive, you need a secure OS and a secure Web server. The Apache server team keeps a close eye on the Bugtraq list and publishes patches. The OpenBSD people are very security conscious. And because their base code is derived from 4BSD, it has been worked over by many of the best developers in the UNIX and inter-networking communities for years. Then, in 1997, OpenBSD itself was given a thorough going over for security, from head to toe. Oliver Friedrichs, vice president of engineering at Secure Networks Inc., Calgary, Canada, wrote to me in January 1998:

"We have been using OpenBSD in our development of security software for the past one-and-a-half years. We develop a network security assessment tool which identifies weaknesses in computer networks. Our main interest in using OpenBSD was its security benefits. We have assisted the OpenBSD development team in a fairly large capacity to find and solve many security problems which were present in OpenBSD and other BSD operating systems. That being said, it's quite some time later now, and in my opinion, OpenBSD has been 'victim' to one of the most stringent software source code reviews (that is, looking for and discovering security flaws), that I know of, occurring in any operating system. Many individuals worldwide have spent much of the last year piling through source code, identifying and fixing security problems.


"... if you want a stable and secure environment, there are definitely many benefits to using it. Most of the security problems found in other free operating systems were fixed first in OpenBSD. For reference, many of these vulnerabilities apply to other operating systems as well, and have been documented as security advisories at SecNet.com/advisories [see <http://www.secnet.com/advisories>]."

The result is that, according to OpenBSD Project Leader Theo Deraadt, not a single remote root security problem was detected between the release of 2.1 and 2.2 (2.3 is scheduled for release this month). For more information on

Benchmark	X86 Running Solaris 2.6	X86 Running OpenBSD
Arithmetic	6.6	7.6
Drhystone 2 (nonregister)	9.0	10.5
Execl throughput	5.4	32.9
File copy (30 seconds)	3.4	5.0
Pipe-based context switch	7.1	13.9
Shell scripts (8 concurrent)	2.2	9.2
Average	5.6	13.2

Figure 2. Benchmark results (comparable to other runs of Byte Benchmark 3.0).

OpenBSD's security policies, see the OpenBSD security page at <http://www.openbsd.org/advisories>. Friedrichs' company produces some exciting security software that is based on OpenBSD. This, and many others, are listed in the OpenBSD software users list at the OpenBSD Web site, which includes companies using OpenBSD internally, many OpenBSD-based ISPs, companies developing or offering software for OpenBSD and so on.



OpenBSD also features a range of cryptographic software, including Kerberos and IPSec, compiled and installed as part of the base system. One could even say the main focus of OpenBSD development is on operating system security. It is, according to the OpenBSD Project, "the only significant OS shipping with full crypto support included." Because OpenBSD ships from Canada, it is only indirectly subject to the U.S. Commerce Department's vagaries and can ship with features that you may never see from SunSoft Inc., The Santa Cruz Operation Inc. or Microsoft.

Even password security is expanded beyond most conventional UNIX varieties. Not only is the shadow password format mandatory, but there are four algorithms for encoding login passwords, one of which is compatible with standard UNIX and the rest of which are more secure. Normally this is not much of an issue because a machine dedicated to Web serving won't have many interactive users. But being on

the Internet, you want your machine as secure as possible. In addition, Kerberos IV authentication is there to allow secure network logins, as for updating files on the Web server.

And network security has not been overlooked. The IPSec protocols ESP and AH are preinstalled on the system, and a pair of tools (a manual configuration process and a daemon for providing standard mappings) lets you control what type of security is used and on which connections. This is a very powerful feature for use on extranets. Overall, OpenBSD security is arguably the best available.

As for reliability, it generally just keeps running. Unlike Windows 3.1, which needs rebooting every hour or so, or Windows 95, which needs rebooting daily, Berkeley UNIX systems tend to stay up for months at a time.

In the early days of 4.2BSD, a bug in the "uptime" code was detected when a system had been up for six months without a reboot and a small integer variable wrapped around. We have not kept our test system up for six months (otherwise, you wouldn't see this review until the end of the year) but, apart from flaky or unsupported hardware, we have never had a system crash or lock up in normal operation.

And in case you do have problems, you can build a kernel with a debugger built in; instead of crashing, system failures will drop it into a debugger. In short, OpenBSD is reliable and secure enough for use as a Web server platform.

Testing Its Speed

We ran the "Byte UNIX Benchmark, Version 3.1" against the same machine running Solaris and OpenBSD. The machine is a Pentium (Cyrix) 166 with 32 MB of RAM (see Figure 2).

The Byte UNIX Benchmark measures CPU time (where the two systems did about equal), the throughput of execing another program (where OpenBSD won hands-down) and other features. The numbers are the ratio of the machine under test to some original, so higher numbers are better.

While benchmarking Web server performance is a trickier proposition, we think these results indicate that OpenBSD is faster than Solaris at many of the things that a Web server does. This should extrapolate to peppy Web server response from OpenBSD, a measure confirmed by the number of ISPs running OpenBSD systems.

Hardware Support

Because OpenBSD is mostly a volunteer effort, the latest hardware may not be supported until somebody gets around to writing support for it. For example, my Dell notebook has a PCI-based PCMCIA bridge, but as of Version 2.2 (and for several months

thereafter), this adapter was not supported (the reason I had to try out the "install from local disk" option when installing on the notebook). As of this writing, we are still using a work-around—telling the system that the PCI-PCMCIA bridge is on the ISA bus—on my notebook. It works just fine but it's a "bit of a hack."

In addition, there is less likely to be support for proprietary hardware. To mention my Dell notebook one last time, it uses some contemptible proprietary "NeoMagic" video adapter, which works just fine under Windows 95 and Solaris 2.6, because the big companies involved play together. But NeoMagic will not give any of its undocumented programming information to organizations like XFree86, so it doesn't support it. And it doesn't appear to emulate anything other than 640-by-480-by-16 VGA mode unless you have the secret documentation, so that's what I use on my notebook. On the deskside PC, I have an ATI video board—ATI Technologies Inc.

seems always to have played fairly with developers large and small, and the ATI boards are well supported at 1,024 by 768 and higher.

On the SPARC, the normal video configurations (CG2, CG3, CG6, BW2 and P4) are supported, as are all of the other standard hardware. There is something to be said for a platform that doesn't have 1,200 different and incompatible video boards available, though cost isn't it. For the latest word on hardware support, check the OpenBSD Web site. There is no MP support at present. This might be an issue of scalability on larger servers.

OpenBSD is a free UNIX-like system that runs on many platforms. It has Perl 5, gcc and many other tools bundled, and much freeware, including the Apache Web server, can be added easily. OpenBSD's meticulous and proactive approach to security means that your system is more likely to survive attacks via the Internet. And the price is very right, too. -->

Try Out The Latest Software



FREE

Software Solutions For The Solaris™ and Java™ Environments.
Sun Solutions NOW!

An All-NEW Web Site From Sun Microsystems

Try It NOW!

Visit the Sun Solutions NOW! Web Site to find new and exciting solutions for the Solaris™ and Java™ environments by Sun's third-party software vendors. Check out the latest applications and stay current on new technologies as they hit the market—sometimes even *before* they're in full release.

Over 500 Software Products

This is the largest collection of downloadable demos and trial software written for Solaris and Java on the Internet. We've added the complete Sun Solutions Catalog with over 14,000 product listings.

Sun Solutions NOW!
www.sun.com/downloads



OpenBSD

Company

The OpenBSD Project
1933 6th St., S.E.
Calgary, Alberta
Canada T2G 2Y3

WWW

<http://www.openbsd.org>

Best Features

A good server OS. Runs on many platforms. Runs many popular Web servers, and Perl for CGI is preinstalled. Freely available in source code, reasonably priced on CD-ROM. Cryptographic and security software included.

Worst Features

Limited support for the wide range of peripherals available on the PC, compared with some commercial operating systems.

Price
\$30

Circle 159

internet expo

*Are you
refocusing on
eBusiness?*

We are.

Internet Expo Presents...
eBusiness World



*The Integration of
Business Vision
and
Technical Implementation
for Maximum ROI.*

Sept. 1-3, 1998 • Boston
Sept. 22-24, 1998 • Toronto
Nov. 3-5, 1998 • Chicago

For a **FREE** Subscription
to *eBusiness News*, visit
www.dci.com/eBusiness

**Reserve Your Exhibit
Space Today! Co-Sponsorship
Opportunities Available.**

produced by



XVGZX00

Circle No. 33

NEW PRODUCTS

The product descriptions are compiled from data supplied by the vendors. To contact them for more detailed information, circle the appropriate reader service number on the card located elsewhere in this issue.

New SPARClike Unveiled

Tatung Science & Technology has added the COMPstation U60-2300, a dual-processor Sun Microsystems Inc. Ultra 60-compatible workstation/server, to its PCI-based product line.

The new COMPstation U60-2300 system comes with two 300-MHz, 64-bit UltraSPARC-II processors that are each equipped with 2 MB of level 2 cache. It supports three 33-MHz, 32-bit PCI devices and one 66-MHz, 64-bit PCI device, and comes with an additional 64-bit UPA slot designed to support add-ons such as Creator graphics cards.

Tatung says the new COMPstation is designed for geophysicists, CAD designers, animators and other high-end graphics users that perform 3D modeling and full digital prototyping.

The U60-2300's five disk drive bays can be fitted with two 3.5-inch hard drives; one 5.25-inch CD-ROM drive, one 4- or 8-mm tape drive; and an additional 5.25-inch slot for another peripheral or a 3.25-inch floppy drive. Other features include Ultra Wide SCSI, two serial ports, one parallel port and 100BaseT Ethernet.

Pricing starts at \$15,010 for a COMPstation U60-2300 with two 300MHz UltraSPARC-II processors, 2 MB of external cache, 128 MB of RAM (expandable to 4-GB) and a 20-inch monitor. It comes preinstalled with Solaris 2.6.

Tatung Science & Technology Inc.
1840 McCarthy Blvd.
Milpitas, CA 95035
<http://www.tsti.com>
Circle 101

High-Capacity RAID Array

Winchester Systems has introduced high-capacity versions of its FlashDisk RAID disk array products, featuring 18.2-GB Ultra SCSI disk drives. FlashDisk offers up to two Ultra SCSI host ports, which deliver 40-MB/s sustained data throughput and maximum rate of 80 MB/s, the company says. The two Ultra-SCSI ports can be connected to a single host for increased throughput, or to two hosts for increased fault tolerance or to share the array hardware.

The 18.2-GB drives are packaged in the company's "hot-swap" drive carriers, which are equipped for Ultra SCSI drives with 68-pin SCA connectors that support the new, high-speed 40-MB/s data transfer rate.


Winchester offers a 145.6-GB work-

Multiservice Access Concentrator

With PortMaster 4 from Lucent Technologies, network service providers can offer a wide range of access applications from a single chassis, including traditional Internet access and new services such as, outsourcing and virtual private networks (VPNs), the company says. For large corporations, PortMaster 4 combines a high-capacity remote access server with remote local-area network (LAN) access and high-performance extranet routing. The product also enables providers to define classes of service—such as email or browser—on a per-use basis, Lucent says.

In a single 10-slot chassis, the PortMaster 4 supports up to 864 simultaneous ISDN or 56-Kb/s modem (V.90 or K56flex) connections.

PortMaster 4 features an embedded 5-GB/s Asynchronous Transfer Mode (ATM) cell-switching fabric and a distributed multiprocessor architecture in which PPP processing, filtering, compression, encryption and packet forwarding are distributed to each card. The switch fabric delivers 155-MB/s full-duplex capacity to each slot—most products on the market deliver 155 MB/s for the entire backplane, Lucent says.



To deliver reliability, PortMaster 4 features a fully redundant, fault-tolerant design, including redundant controller, switch fabric, power and fans, Lucent says.

The product runs on Lucent's ComOS operating system, which includes support for IP routing protocols such as RIP, Open Shortest Path First (OSPF) and Border Gateway Protocol Version 4. PortMaster 4 supports both SNMP management as well as Lucent's Java-based PMconsole configuration

and management utility, and will be supported in a future version of Lucent's OneVision network management platform, the company says.

List price for a fully configured system with chassis, power supply, system manager module with 100-MB/s Ethernet, modem modules and configuration, management, routing and WAN software is \$519 per port.


Lucent Technologies Inc.
4464 Willow Road
Pleasanton, CA 94588
<http://www.lucent.com>
Circle 100

group pedestal unit with two Ultra SCSI host ports and eight 18.2-GB disk drives for \$37,826, and a 200.2-GB workgroup pedestal unit with 12 18.2-GB drives for \$50,355. The FlashDisk Ultra SCSI host ports support Solaris, AIX, Digital UNIX, Open VMS and HP-UX.

Winchester Systems Inc.
400 West Cummings Park
Woburn, MA 01801
<http://www.winsys.com>
Circle 102

Fiber-Optic Modem with Universal Interface

Telebyte Technology has introduced Model 8277, a fiber-optic modem with a universal interface. Model 8277 is designed with switchable interfaces for RS-232, RS-422 and RS-485. The RS-485 interface can also be configured to



operate either half or full duplex on two or four wires. RS-485 options can also enable the fiber transmitter from RTS or data. The fiber port uses ST connectors to accommodate multimode fiber and operates at a wavelength of 850 nm. The modem is designed for use with any application using the Deutsche Industrie Norm (DNI) rail concept.

Network links can be established that can handle data transfer rates of up to 1 MB/s. In addition, Model 8277 has a power budget of 12 decibels that provides reliable operation over fiber distances exceeding 2km. It comes in a 3-by-4.5-inch DIN rail enclosure and costs \$249.

Telebyte Technology Inc.
270 Pulaski Road
Greenlawn, NY 11740
<http://www.telebyteusa.com>
Circle 103

Web Catalog Software on Solaris

Open Market is now offering its LiveCommerce Internet catalog software for Sun Microsystems Inc.'s Solaris operating environment. With LiveCommerce users can develop and maintain commerce-ready catalogs via the Web.

The software reportedly features adaptive navigation, allowing customers to search and view the catalog in a variety of ways, including browsing, hierarchical, text and parametric searching (such as filtering, sorting and drilling down to detail). In addition, the software comes with design templates that can be customized to reflect the business/brand identity.

LiveCommerce includes a suite of Java-based modules for remote administrative access to the catalog database and integrates with existing systems via a real-time query API. It also includes built-in industrial application logic, a robust parametric search engine and a library of more than 50 pages of display objects.

Pricing for LiveCommerce for Solaris with one development license, one production license and an unlimited number of authoring seats begins at \$45,000.

Open Market Inc.
1 Wayside Road
Burlington, MA 01803
<http://www.openmarket.com>
Circle 104

Reverse Engineering of Java Code

Aonix has announced the release of Software through Pictures (StP) Version 2.5, a visual modeling tool set that offers Java developers the ability to generate Java code from Unified Modeling Language (UML) models and supports reverse engineering of existing Java code.

StP supports the Object Management Group UML 1.1 standard, as well as the Object Modeling Technique (OMT), Booch and Structured and Data Modeling, and is targeted at organizations committed to designing large, complex Java applications. In addition, developers are able to reverse-engineer existing Java code and represent that code in the form of visual models. These visual models can be expressed in industry standards notation such as UML. Once captured as visual models, users can modify the application

using StP and regenerate the Java code for the final application.

Version 2.5 now features support for CORBA IDL and full integration with Quality Systems and Software Inc.'s DOORS, a tool that allows users to capture and manage a product's requirements. It runs on Solaris 2.6, HP-UX 10.2, AIX 4.1 and Windows NT 4.0. Pricing starts at \$4,500 for the NT version and at \$8,300 for the UNIX version.

Aonix Corp.
595 Market St., 10th Floor
San Francisco, CA 94105
<http://www.aonix.com>
Circle 105

Java-Based Communication Suite

Skunk Technologies has introduced the Entao Extended Enterprise Suite (EES), a 100% Pure Java suite of applications for enterprises with multiple platforms and/or geographically dispersed locations. The suite comprises five applications: Entao Email, a POP3/IMAP4-compliant email system; Entao Text Chat, for group or one-on-one discussions over a network; Entao Visual Chat, for auditorium-style presentations or communications requiring graphics capabilities such as graphs, charts or diagrams; Entao Forums, threaded discussion groups with support for encryption, digital signatures and security log-in protection against unauthorized access; and Entao Contact Manager, for keeping track of customer and/or personnel information.

The Entao EES is designed to be a network-centric solution for collaborative computing and communication among multiple platforms, including PCs, Network Computers, workstations and any other device equipped with a Java Virtual Machine (JVM). The Entao applications support open standards such as CORBA/IIOP, JavaBeans and Sun Microsystems Inc.'s 100% Pure Java specification.

The software is installed on a host machine, such as a UNIX or Microsoft Corp. Windows NT server, running a JVM. All configuration management, real-time monitoring and event notification is done on the host. The applications can be accessed across the network by any client with a JVM.

Entao EES is priced at \$20,000 per server for an unlimited number of users.

Skunk Technologies Inc.
1437 7th St., Ste. 300
Santa Monica, CA 90401
<http://www.skunktech.com>
Circle 106

Middleware Delivers Legacy Data to the Web

Hitachi Computer Products has introduced TPBroker 3.0, transaction processing middleware for the Object Request Broker (ORB) market. TPBroker 3.0 integrates transactional technology from the database and mainframe worlds using object-oriented technology.

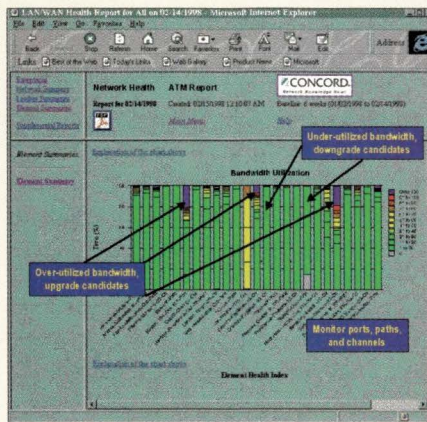
According to Hitachi, TPBroker 3.0 is specifically tuned for Inter/intranet transactions and fully supports the IIOP protocol and the Open Management Group's Object Transaction Service (OTS) 1.1 specification. It supplies connectivity to commercially available relational databases such as those from Oracle Corp., Informix Systems Inc. and Sybase Inc.

TPBroker's key transaction processing features include flat and nested transactions for complex applications and design-level compatibility with existing systems; explicit propagation and direct/indirect context management for flexible application design; internal component architecture for integration flexibility, administration and management; and multithreaded OTS. Pricing for TPBroker 3.0 starts at \$2,745 (development licenses) and \$4,165 (deployment). Evaluation copies are available via <http://www.tpbroker.com>.

Hitachi Computer Products Inc.
2000 Sierra Point Pkwy., MS 500
Brisbane, CA 94005
<http://www.hitachi.com>
Circle 107

Reporting Tool for ATM Networks

Concord Communications has introduced Network Health-ATM, a Web-based reporting and analysis tool for Asynchronous Transfer Mode (ATM) networks. Network Health-ATM interprets data from a variety of ATM vendors, including Bay Networks Inc., Cisco Systems Inc., FORE Systems Inc., Nortel (Northern Telecom) and NetEdge Sys-



tems Inc. It aggregates ATM traffic patterns, establishing baselines across more than 10,000 ports, paths and channels to provide a single, correlated view of key ATM performance statistics. ATM managers can now view all critical usage and performance statistics from the physical port, logical path and logical channel layers. In addition, it is said to allow organizations to balance network resources and make better use of expensive bandwidth.

Because Network Health reports are Web-based, network managers can monitor network performance from any browser and access can be password protected. Available in Solaris, HP-UX or Windows NT versions, Network Health-ATM is targeted at large companies, Internet service providers or telecommunications companies and costs \$3,000.

Concord Communications Inc.
33 Boston Post Road W.
Marlboro, MA 01752
<http://www.concord.com>
Circle 108

Notes/Domino Backup Unveiled

Legato Systems has introduced NetWorker BusinessSuite Module for Lotus Notes, a product that offers online, nondisruptive backup for both Lotus Development Corp. Notes and Domino environments running Solaris, AIX and Windows NT, Legato says.

The module provides Notes customers with a storage management system that automates the backup process, increases data availability and provides disaster recovery support, Legato says. It performs full and incremental backups online and offers the ability to back up all


critical Notes files and create a secondary copy for off-site storage. It can also integrate application backups with file system backups for complete enterprise storage management, Legato says.

NetWorker BusinessSuite Module for Lotus Notes costs \$6,000 per Notes/Domino server operating on UNIX platforms, and costs \$1,200 per server on Windows NT.

Legato Systems Inc.
3210 Porter Drive
Palo Alto, CA 94304
<http://www.legato.com>
Circle 109

Connect Ethernet to Fiber Backbone

Allied Telesyn has announced a cost-effective solution for interconnecting copper- and fiber-based Ethernet systems without the expense or complication of an external fiber transceiver. The AT-MR840TF is a hub with eight 10BaseT ports and one 10BaseFL fiber uplink port.



As many as four AT-MR840TF hubs can be cascaded to form a workgroup of up to 26 nodes. The cascade is accomplished through a separate 10BaseT uplink port equipped with an MDI crossover switch. The hub can be attached to the fiber backbone via either an SC or ST fiber connector. Each hub port is capable of detecting and correcting wiring crossovers as well as auto-partitioning after the detection of repeated collision errors.

Diagnostic LEDs for each port simplify installation, fault isolation and error detection for easy maintenance, Allied Telesyn says. The AT-MR840TF is priced at \$299 and comes with a lifetime warranty.

Allied Telesyn International Corp.
19015 N. Creek Pkwy., Ste. 200
Bothell, WA 98011
<http://www.alliedtelesyn.com>
Circle 110

The Server/Workstation Marketplace

The high impact, cost-effective and powerful advertising section reaching over 93,000 technical product specifiers and buyers.

To advertise call Carol Flanagan at
(508) 839-4016

Email: caflanag@aol.com

And for information on the
UNIX/NT
Recruitment Section
 call the above

SUN-IBM-HP

SOFTWARE	LIST	OURS
WordPerfect 7.0 1st user	495	395
Netscape SuiteSpot Pro	6995	6145
VSI*FAX Gold Fax Software	2395	1995
VSI*FAX Gold Windows Client 10-user	590	520
FrameMaker 1-user (floating license)	2500	2295
Adobe Acrobat 1-user	295	245
Informix On-Line Dynamic Server 1-usr (min. 10)	1500	call
Informix On-Line WorkGroup Serv 1-usr (min. 5)	295	call
Informix 4GL Development 1-user (min 3)	900	call
Term Communications Software	695	595
LP Plus Print Manager/Spooler 4-printer	695	575
SoftWindows MS-Windows Emulator	549	475
Spatch alphanumeric paging SW 4-user	499	425
DoubleVision Remote Support S/W	795	665
Uniplex II+ (20-user)	5500	4695
Disk_Pak Disk Optimizer/Delragger	1495	1230
Cheyenne Arcserve/Open Backup S/W	1995	1395
Legato Networker Backup S/W	2500	1950
Hummingbird PC X-Server Win NT/95	545	395
Hummingbird Maestro NFS Server - Win NT	195	169

HARDWARE	LIST	OURS
HP LaserJet 5Sinx 24ppm, 4MB, ENET	3449	2895
Mannesman Tally MT-661 800pm Line Printer	7999	6795
HP 8GB External 4mm DAT Tape Drive	1370	1175
HP 24GB External 4mm DAT Tape Drive	1819	1495
HP 40GB External DLT Tape Drive	5723	4925
Exabyte Mammoth Ext. 40GB DLT Tape Drv	5295	3995
Tandberg 2.5GB 1/4" External SCSI Tape Drv	969	870
Seagate Cheetah 9, 1GB Ext. Ultra SCSI HD	1995	1375
Plextor 20X External CD-ROM Drive	395	255
Young Minds CD-Studio (ntwk CD-writer Sys.)	8339	Call
DigiPortServer II 16-port Terminal Server	2395	1890
Digi 16 port Serial I/O card/concentrator	2395	1890
Central Data SCSI Terminal Server - 16 ports	1295	1095
Multitech 33.6 External FAX/DATA Modem	269	145

Other Platforms: SCO, SUN X86, WIN NT, DEC, SGI, DG

Open Systems EXPRESS!

Your direct source for UNIX,
 Win NT & networking products
 45 Whitney Road, Suite B5, Mahwah, NJ 07430

100's of other UNIX products available!

Free Shipping! **1-800-445-9292** Fast Delivery!
<http://www.osexpress.com>

(Or call and ask for our free catalog!)

Fax: 201-847-8899 MCVISA/AMEX/
 E-mail: wiz@osexpress.com DISCV/C.O.D./Net.

Circle No. 353

Memory Upgrades

Significant Savings • 100% Compatible
 Lifetime Warranty

Complete line of
 SUN, SGI, IBM and NT
 Memory Upgrades



Now shipping!

1GB Sun 450 oem pn: x-7005a
 1GB Sun Enterprise 3000-6000 oem pn: x-7023a
 512MB SGI Octane oem pn: hu-mem-512
 256MB SGI O2 oem pn: mem-256-02

Order on-line fast, easy, secure at
www.mcac.com
 USA toll-free: 1-888-MCA-7898
 email: sales@mcac.com

FREE worldwide
 delivery on all orders
 over \$5,000!



MCA Computer Corporation

Authorized Distributor



The Gigabyte Memory Specialist

Circle No. 372

MICHAURA SYSTEMS CORPORATION

...the upgrade specialists

Buy... Sell... Trade

Sun Microsystems
 and more...

- Systems
- Spares
- Lanotronix
- Best Power
- Repairs
- Advance Replacements
- Peripherals



Tel: 781-937-0010
 Fax: 781-937-0808
 Email: michaura@aol.com

All trademarks are the property of their respective owners.

Circle No. 330

When It Comes To Computers...



We Make The Sun Rise Again

800-566-4SUN

Sell-Buy-Rent-Exchange



"A Family Business"

101 First Street Utica, NY 13501

Phone: (315) 724-2209 FAX: (315) 724-0794

www.ccnny.com

Sun Workstations/Servers
IBM RS6000 Business Partner
Sun 20" Monitor Repair
Sun Parts & Peripherals
SONY Authorized Dealer
Netscape Solution Experts



Circle No. 420

JUST OFF LEASE

SPARC 5/85
\$1,950. Complete

This is the lowest price SPARC 5/85 SYSTEM we have ever offered

System includes:
SPARC 5 Model 85 • 32Mb Mem
1.1Gb Internal Disk • TX Graphics
16" color monitor - keybd kit

Also several SPARC 5/70 (no monitor)
\$1,475.

Limited supply, call now!



800-456-6233
FAX (714) 632-9248

Circle No. 324

SYSTEMS & CPU's

S20TX1N-71-32-P46	\$ 4,975.00
S20TX1N-50-32-P46	\$ 3,475.00
U1/170-64mb-2.1GB-CD-FL	\$ 4,975.00
U2/2300-512mb-8.4GB-CD-FL	\$18,700.00
4/15CN-32-P43	\$ 775.00
S5TX1N-170-32-P46	\$ 2,675.00
SM81 85mhz	\$ 2,300.00
SM71 75mhz	\$ 1,650.00
HS151 150mhz	\$ 2,100.00
U2/167mhz	\$ 975.00
U2/200mhz	\$ 1,450.00
U2/300mhz	\$ 3,950.00
S5/170 cpu/sys brd	\$ 1,600.00

MEMORY & DISKS

1GB SS20/S5/Ultra	\$ 175.00
2GB SS20/S5/Ultra	\$ 350.00
23GB External	\$ 2,100.00
X164PC 64mb SS20/Ultra	\$ 350.00
X132PC 32mb SS20/Ultra	\$ 225.00
X132MC 32mb SS5	\$ 175.00

GRAPHICS

TX1/TX4	\$ 375/775
CG3/GX	\$ 165/275

ENTERPRISE

EX000 X2600A System board	\$ 5,200.00
EX000 X2601A System board	\$ 6,700.00
EX000 X2610A I/O board	\$ 4,750.00
EX000 X7022A 256mb	\$ 1,575.00
EX000 X7023A 1GB mem	\$ 4,950.00

SPECIAL 9GBEXT 9GB external Seagate
\$ 675.00

Large inventory of hardware including spare parts, workstations, peripherals, memory, and 3rd party products.

12 month warranty with same day shipping on all in stock items.

Check our web page for weekly specials

Email: sales@gshiis.com

techinfo@gshiis.com

www.gshiis.com

* Call for full product line and availability. * Trades accepted, and purchase options available on all excess used equipment.



GSH Intelligent Integrated Systems, Inc.

95 Fairmount Street Fitchburg, MA 01420

Tel: (603) 529-7880 Fax: (603) 529-7884

Circle No. 413

To Advertise Or For More Information Concerning The

UNIX/NT
Recruitment
Opportunities Section

Please call Carol Flanagan at
(508) 839-4016

RENT
THIS
SPACE

Call Carol Flanagan
(508) 839-4016



SUN MICROSYSTEMS
WORKSTATIONS
PRINTERS SERVERS MEMORY
OPTIONS DISK TAPE

1-800-444-7003

BUY • SELL • RENT • LEASE •
UPGRADE • EXCHANGE

InterContinental
Computers, Inc.

4824 W.96th St., Minneapolis, MN 55437

VOICE (612) 835-4555 FAX (612) 835-3936

EMAIL: 1grif@worldnet.att.net


Circle No. 317



It's Not Just for the Desktop Anymore.

Liberate compression from the desktop, and you've got a strategic business weapon. Such vision prompted the initial release of PKZIP UNIX in October 1992, signalling the advent of open systems compression.

As the pioneers in the field, we've developed compatible compression for 34 new operating systems to revolutionize the way you do business. Call us today or visit us on the Internet for more information.



PKZIP UNIX Operating Systems

AIX • AT&T 55V4 • BSD 386 • Coherent • DEC UNIX
 DYNIX PTX • 88 OPEN/IG AVIION • HPUX
 ICL DRS/6000 • LINUX • MIPS 1 • MIPS 2
 MIPS/SGI • Novell UnixWare • SCO • Solaris x86
 Solaris/SUN OS • ULTRIX • Univel • XENIX

- 100% cross-compatible
- Built-in data encryption
- One step multiple file archiving and compression (no \$ tar)
- 32-bit CRC error checking
- Directory recursion
- Reduce file size and data transmission costs by 90%




Download
 YES IT NOW!
www.asizip.com/Sun



Phone: (937) 847-2374 · Fax: (937) 847-2375 · E-mail: sales@asizip.com

Circle No. 376



**Customer Blown Away
 By Low Prices**

Work Stations • Servers • Peripherals • Parts
 Sales • Rentals • Leases • Exchanges • Repairs
 Maintenance • 120 Day Warranty

Call today to buy, sell or trade SUN and
 Silicon Graphics equipment with Security!



622 Rossmor Building • 500 North Robert Street • St. Paul, MN 55101
 612/227-5683 • FAX: 612/223-5524 • seccomp@seccomp.com

Circle No. 314

**NIPrint™ Completes
 Win/95/NT-to-Unix Integration**

NIPrint™ is both a LPR and a LPD print server for Winsock. NIPrint is installed on the Windows PC only. NIPrint works with Unix's LPD/LPR or your print servers LPD - no host files or modification required. NIPrint is fully compatible with Windows, Windows 95 and Windows NT (3.51 and 4.0).

- BI-DIRECTIONAL, WORKS WITH LPD ENABLED PRINTERS/DEVICES
- IMMEDIATE PRINTING (NO WAITING, NO POLLING)
- PRINT JOBS ARE TRANSPARENTLY SPOOLED TO/FROM UNIX
- FULLY LPR/LPD COMPLIANT - WORKS WITH ANY UNIX OR PRINT SERVER (EMULEX, JETDIRECT, TEKTRONIX, ETC)
- FULLY INTEGRATED INTO PRINT MANAGER, JUST CLICK FILE, PRINT
- COMPLETE 16- AND 32-BIT VERSIONS INCLUDED
- CAN BE RUN AS A SERVICE IN NT OR NTAS (3.51 AND 4.0)
- INSTALLS IN MINUTES - REALLY!

DEMO ALSO AVAILABLE FROM OUR WEB SITE
 OR CALL (612) 932-9899 FAX (612) 932-9545



www.networkinstruments.com


**NIPRINT STARTS AT
 \$29.00 PER PC**

5-PC \$119, 10-PC \$189
 25-PC \$349, 50-PC \$595
 HIGHER USER COUNTS AND
 UNLIMITED SITE LICENSES AVAILABLE

CALL 800-526-7919 FOR A FREE 21-DAY DEMO

Circle No. 374

ULTRAFast SBus A/D



ALSO PCI!

10 MHz DATA ACQUISITION

The ultimate 12-bit data acquisition board can acquire 8 Million A/D samples while simultaneously outputting 8 Million D/A samples and TTL vectors, all without CPU intervention.

Four analog inputs, twin 5 MHz 12-Bit A/D and D/A converters, 5 MHz TTL input/output vectors, and huge 16 MB RAM allow use in the most demanding scientific and industrial applications.

ULTRAVIEW CORPORATION email: ultraview@worldnet.att.net
 34 Canyon View, Orinda, CA 94563 www.ultraviewcorp.com
 Telephone: (925) 253-2960 Fax: (925) 253-4894

Circle No. 375

PS. Need a hard-to-find storage device?



Whether you need leading edge technology or a tried and true product, Peripheral Solutions specializes in storage and memory for your SUN, SGI or NT system. We support everything we sell, and our in-house technical phone support, repair technicians and 24 hour advance exchange program will get you back on-line fast.

Our sales team will help solve your storage needs with disk drives, tape drives, CD-R, memory, media, workstations and RAID's. We offer a complete line of storage and memory solutions, and we only carry the top manufacturers, including new and out-of-production products from:

- Seagate – Hard Drives, Tape Drives
- Kingston – Memory
- Quantum – Hard Drives, DLT
- Sony – CD-R, Tape Drives
- IBM – Hard Drives
- Tandberg – Tape Drives
- Exabyte – Tape Drives
- CMD – RAID Controllers

PS. Relax, and let Peripheral Solutions work for you.

PERIPHERAL SOLUTIONS 

1-800-255-3475 Fax 408-425-5792

www.alldisk.com

Circle No. 318

The Server/Workstation Marketplace



The ultimate in Windows to UNIX connectivity

FacetWin makes Windows®95/NT* to UNIX integration easy and affordable.

Just ask for a **FREE EVALUATION** copy of FacetWin, or check out our web site at www.facetcorp.com.

• **File & Print Services**

No need for NFS or ftp... files and printers on UNIX systems simply appear as local resources to your PC, and are ready to use!



• **Terminal Emulation**

World class terminal emulation from the terminal experts that brought you FacetTerm*. Set up a UNIX application with a Windows icon in one minute!

• **Plus, check out these cool features:**

- PC Backup/Restore
- Modem Pool Server
- Remote Computing Support
- Email Server

Connecting Windows to UNIX... the Windows Way!

FacetCorp 

tel: 800/235-9901 • 972/985-9901

fax: 972/612-2035

info@facetcorp.com

www.facetcorp.com

FacetWin is a registered trademark of FacetCorp. Other names are properties of their respective holders.

Circle No. 378

SUNEXPERT REPRINTS

Keep up with the latest in Sun technology with the best minds in the industry. Use reprints to promote, inform, and sell.

Reprint Management Services™ provides you the opportunity to obtain reprints of reviews, articles, and features in **SunExpert Magazine**.

High-quality editorial reprints will help your company in many ways:

- Increased EXPOSURE for your product or service
- Credible, believable information that consumers TRUST
- Excellent SALES tool for trade shows, mailings and media kits
- Powerful EDUCATIONAL RESOURCE for consumers and employees

Reprints are completely customized to your needs. Call today for additional information!

CATHY HENSEL

Reprint Operations Specialist

REPRINT MANAGEMENT SERVICES™


147 West Airport Road, P.O. Box 5363
Lancaster, PA 17606

Phone: (717) 560-2000

Fax: (717) 560-2063

Email: sales@rmsreprints.com

REMOTE YOUR TAPE BACKUP JUST ABOUT ANYWHERE



POWERLINK connects your Ultra wide SCSI devices to just about anywhere

- Extends the SCSI bus to 10Km
- Ultra / wide 40 Mbytes/sec.
- Single-ended and differential versions
- Real time access to remote SCSI peripherals

Call today for more information about
POWERLINK, and other SCSI connectivity
solutions from APCON.

1-800-624-6808 x 207




APCON, Inc.
17938 SW Upper Boones Ferry Rd.
Portland, OR 97224
Tel: 503-639-6700
Fax: 503-639-6740
E-mail: sales@apcon.com



Visit our WEB site at www.apcon.com

Circle No. 442



Customer's Head Spins Over Great Deals

Call today to buy, sell or trade SUN and Silicon Graphics
equipment with Security! Unbeatable value, unparalleled
service -- no wonder so many heads are spinning!



622 Rossmor Building • 500 North Robert Street • St. Paul, MN 55101
612/227-5683 • FAX: 612/223-5524 • seccomp@seccomp.com

Circle No. 401



E.L.I. SYSTEMS inc.

Used SUN Equipment

139 Hampshire St., Cambridge, MA 02139

Design Your Own System or Purchase Parts Individually

Start with a SUN BASE

Ultra 1/140/170	\$3,700/4,000
Ultra 2	7,200
Sparc 20	1,200
Sparc 10	1,100
Sparc 5 70/85	900/1,000
Sparc 5 110/170	1,300/2,100
4/75B Sparc 2	235

Add a MONITOR

16" Sony color	\$275
17" Sony color	675
19" Sony color	475
20" Sony color	975

Choose a CPU

200/300 Mhz, Ultra 2	\$1,200/3,800
501-2352/2562SM51	450
501-2752/2769 SM61	595
501-2925 SM71	1,795
Ross 125 Mhz/512K	new 1,925
Ross 150 Mhz/512K	new 2,495
Ross 180 Mhz/512K	new 3,120
Ross 200 Mhz/512K	new 3,500

Select a GRAPHICS CARD

501-1845 ZX F/B	\$595
501-2253 TX4 F/B	695
501-2325/2922 TX1 F/B	400
501-1645 D/S GX F/B	175

Remember MEMORY

X164P 64 MB SP 20/10	\$450
X164F 64 MB SP10	400
X132P 32 MB SP 20	175
X132M 32 MB SP5	225
X116F 16 MB SP10	65
501-1739 4MB S2/1/IPC	15
3 rd - Sparc 5, 32 MB	130
3 rd - Sparc 10/20/Ultra, 32 MB	145
3 rd - Sparc 10/20/Ultra, 64MB	225

What capacity DISK

Seagate 1GB N/W/C	\$175/295
Seagate 2.1 GB N/W/C	395/435
Seagate 4.2 GB N/W/C	650/675
Seagate 9.0 GB	925

SYSTEMS SPECIALS


A11-UAA1-1A-96AB 19" mon	\$5,900
S20TX-51-32-P46 19"cm	3,300
S10FGX-51-64-P46 16"cm	2,750
S5FX1-110-32-p46 17"cm	3,095
4/75FGX-32-P46 16"cm	1,095

PHONE 1-800/447-1156
FAX 617/354-1417

E-MAIL sunsales@eli.com
WEB <http://www.eli.com>

Circle No. 428

IBM * SUN * HP



- WORKSTATIONS • SERVERS
- PERIPHERALS • PARTS

BUY - SELL - LEASE

NEW & USED EQUIPMENT

UNIX Certified Warehouse

800-323-9751

Fax:847-228-0551



www.foj-gti.com
marketing@foj-gti.com

FRANK ORLANDO JR. & CO., INC
501 W. Algonquin Rd.
Arlington Heights, IL 60005

Circle No. 325

Over \$1 Million In Off-Lease Inventory!

Gulfc^oast
Workstation Corp.

2210 Tall Pines Drive, Suite 220
Largo, FL 33771

Phone: (813) 587-7882

Fax: (813) 587-7974



CALL FOR CURRENT PRICING
(800) 587-7882

Circle No. 409



Get in Touch
with Trident.

TOUCH SCREEN TECHNOLOGY EXCELS AT
Simplifying computer input
Saving space
Supporting fast, frequent computer operations
Reducing maintenance costs and down-time

AND NOW, YOU CAN HAVE TOUCH WITHOUT
MODIFYING YOUR EXISTING APPLICATIONS.
TRIDENT'S TOUCH SOLUTIONS INCLUDE
Touch-configured monitors, X terminals and Netstations
Touch screen integration with your monitors
Touch screen hardware and driver software
All major touch technologies
Rack mounting or ruggedized NEMA enclosures
and kiosks with touch
Touch application software development services

TRIDENT'S TOUCH SOLUTIONS WORK WITH
UNIX Workstations X Terminals UNIX PC's
DOS & Windows Netstations Macintosh

Trident puts you in touch.™

TRIDENT

TECHNOLOGY SOLUTIONS

TRIDENT SYSTEMS INC. email: touch@tridsys.com
TEL 703.273.1012 url:http://www.tridsys.com/cpg
FAX 703.273.3763

Circle No. 408

Introducing...



The unified expertise of Amcom Corporation and Delta Parts, Inc.



◆ **RS/6000** ◆
SAVINGS UP TO 75%!

- ◆ **Whole Machines**
- ◆ **Parts**
- ◆ **Repairs**
- ◆ **Logistics Management**
- ◆ **Features**
- ◆ **Upgrades**

Contact Tim Balko
or Jeff Karschnik

800.328.7723

Or Via e-mail
tbalko@amcomcorp.com
jeffk@amcomcorp.com

6205 Bury Drive
Eden Prairie, MN 55346

Circle No. 431



Customer Goes Ape Over Great Service


Call today to buy, sell or trade SUN and Silicon Graphics equipment with Security. Unparalleled service, unbeatable value – No wonder so many people are going ape.



622 Rossmor Building • 500 North Robert Street • St. Paul, MN 55101
612/227-5683 • FAX: 612/223-5524 • seccomp@seccomp.com

Circle No. 438

CD ROM STORAGE



New!
Mohave Series

SUN/Solaris™ Based CD-ROM Intranet/Internet Servers with up to 500+ Disc on-line.

Share Historical/Archive Data Company/World Wide

Featuring:

- ▼ Price-Performance
- ▼ Interoperability
- ▼ Scalability
- ▼ Ease of use
- ▼ Reliability

Based on Sun Microsystems' SPARCengine™ Ultra™ AXI motherboard and Solaris™ software, the Mohave Server allows access to over 300 GB of data via HTTP, NFS, or FTP over a TCP/IP connection. Configurations available with CD-Recorder/Reader combinations.

For the latest in CD-ROM storage, visit our web site: www.boffin.com




Call... **800.248.5328**

2500 West County Road 42, Suite 5
Burnsville, Minnesota 55337
Bus: 612.894.0595 Fax: 612.894.6175
Email: sales@boffin.com

Circle No. 363

A Monochrome Solution In Stock That Adapts to Your SUN System

The MONORAY MR2000 SUN monochrome monitor is specially designed to work with the SUN MG1 monochrome frame buffer. The result is superior performance at all frequencies, not just optimized for one portion of the available range.



Features:

- 20" Flat Profile High Resolution
- Multi-frequency Monochrome
- 3 Year Warranty
- Resolution Up to 2048 x 1536
- ECL Interface
- Analog Interface Available



Distributed by:



Rely on Richardson for superior performance, digital accuracy, and consistent quality at a competitive price.

800-222-CRTS(2787) www.rell.com/crts dpg@rell.com

Circle No. 454

RENT to OWN
SUN
Sparc 20 • 250/mo*
Sparc 20 TX-61-64-P46(2)
 TX1 Graphics
 61MHz Processor
 64Mb Memory
 20" Color Monitor
 2 x 1.05GB Disks
 Keyboard Mouse & Pad

*Lease for 24 mo. @ 250./mo then it's yours for \$1.00 - or - purchase for \$4950. Just off lease from a major auto mfg. Subject to supply on hand. Lease subject to credit approval.


(800) 456-6233
FAX (714) 632-9248

Circle No. 455


To Advertise Or For More
 Information Concerning The
**UNIX/NT Recruitment
 Opportunities Section**
 Please call Carol Flanagan at
(508) 839-4016

Sun & HP
BUY SELL TRADE
Workstations and Servers
Boards Monitors Disks Tapes Mass Storage Solutions

- ◆ Same Day Shipping
- ◆ Competitive Pricing
- ◆ Custom Configurations Available

DEPOT REPAIR CENTER

- ◆ 30% To 50% Cost Savings
- ◆ One Year Warranty
- ◆ Overnight "Swap" Service


ACC, Inc.
AdvanTec Computer Company
 4125 Business Center Drive Fremont, CA 94538
 Phone: (510) 440-9700 Fax: (510) 440-9709
 E-mail: sales@advanteco.com
 http://www.advanteco.com

Circle No. 439

JavNet SUN Microsystems Oracle

One Stop SOLUTION For All Your Networking Needs:
 Sales/Integration/Support/Consulting

(888) DIGICAD

Software/Compilers
 Softwindows95 for Solaris 2.6 \$ Call
 C++ Compilers, Java JDK

Options/Peripherals
 Memory, Disk, Monitor, Printer Upgrades

NetWork: LAN/WAN
 Cisco Routers/Hubs/Switches/Devices

Web Server
 Digidac JavServ Solaris 4/300MHz
 - 4x300MHz, 2GB RAM, 84GB Disk
 \$ Call

Digidac JavServ NT 2/333MHz
 - 2x333MHz, 1GB RAM, 90GB Disk
 \$ Call

Sun® Ultra Entp450
 Sun WorkGroup UltraE450 \$ Call
 - 250MHz, 512 MB RAM, 42 GB Disk,
 12XCD, Server License
 \$ Call

UltraSparcII - 2x300MHz, 512MB
 RAM, 18GB Disk, 12XCD
 \$ Call

Internet Services


- T1 from 56K up to 1.5MB
- Multiple T1 and T3 + ATM
- DS3 and OC3
- Web Design, Hosting, Advertising, etc.
- Web Security & Monitoring Software

Oracle®

Oracle Solution Provider:

- Consulting, Developing, Sales Integration, Support to fit to your Corporate RDBMS needs

\$ Call


DIGICAD CORP.
 Phone (949) 724-8600 * Fax (949) 724-9151
 2819 McGaw * Irvine * California * 92614
<http://www.digidac.com>

Circle No. 352

Why Pay More?

IBM RS/6000

From IBM National complete Line distributor

IBM GUARANTEED

SAVE OVER GOING DIRECT TO IBM

NEW FULL LINE SUPPLIER / ALL SYSTEMS IN STOCK


- Special:
- 3153 - BG3 \$345
 - 7208 - 341 20-40 GB 8MM \$3,900
 - IBM 9.1 GB (7204) F/W Disk Drive \$1,620

call for **free CATALOG**



National Data

EMAIL: RS6000@natdata.com



IBM
• MEMORY •
CPU upgrades

CALL : 800 - 406 1811

Circle No. 343

RS/6000

**NEW • REFURBISHED
BUY • SELL • LEASE
SYSTEMS • PARTS
PERIPHERALS**



Worldwide Trade Corporation

5253 Edina Industrial Blvd.
Edina, MN 55439
Call (612) 831-8094
Fax (612) 831-7018

Circle No. 370

RS/6000 BUY ~ SELL ~ LEASE

- NEW & RECERTIFIED RS/6000 PRODUCT
- IN-HOUSE TECHNICAL SERVICES
- SAME DAY SHIPPING (UNTIL 7:00 PM CST)
- UPGRADES - PARTS - FEATURES
- KINGSTON PRODUCT DISTRIBUTOR
- SHORT-TERM RENTAL/LEASE



EVOLVING

**SPECIALISTS IN RS/6000
PRODUCTS & SERVICES**

SOLUTIONS

1.800.294.4362

E-MAIL: EVOLVING@EVOLVINGSOL.COM
WEB [HTTP://WWW.BUSINESS.NET/EVOLVINGSOL](http://WWW.BUSINESS.NET/EVOLVINGSOL)



6525 CECILIA CIRCLE ▼ BLOOMINGTON, MN 55439 ▼ 612 . 944 . 9200 ▼ FAX: 612 . 944 . 9292

Circle No. 388

RS/6000

New & Used

MSI

RS/6000

New & Used

Minnesota Systems International, Inc.

CALL TODAY 800-947-0811

Systems • Peripherals • Memory

BUY • SELL • IBM • DIGITAL • SUN

Former IBM RS/6000 sales staff available for *immediate* answers to all your hardware questions.

1701 East 79th Street • Suite 11 • Minneapolis, MN 55425

Phone: (612) 883-0808 FAX: (612) 883-0893

msi@bitstream.net



Circle No. 391

**RENT
THIS
SPACE**

Call Carol Flanagan
(508) 839-4016

Lightning FAX 5

Windows & Java Clients
Least-Cost Routing Management
Fax Load Balancing
Native Fax Server on IBM® AIX®
& Windows® NT®

**E-MAIL-TO-FAX GATEWAY
API WITH FEEDBACK
BROADCAST MODULE
FAX CALL CENTER**

**Interstar
TECHNOLOGIES**

Toll-free (888) 766-1668
sales@faxserver.com

www.faxserver.com

Circle No. 327

RECRUITMENT ADVERTISING

Reach over
93,000
IT professionals with a recruitment ad in the
UNIX/NT career opportunities section of
SunExpert magazine. For special recruitment rates call:
Carol Flanagan at
1-508-839-4016

**WebServer
OnLine Magazine**

The latest information to help Web
professionals manage and optimize their
site is just a click away.

<http://webserver.cpg.com>
NO PASSWORD REQUIRED

Discover Solar Systems with hundreds of Suns.

Refurbished Sun Microsystems equipment at great prices. Fully-tested inventory on hand, same-day shipping, satisfaction guaranteed.



- Workstations & Servers
- Upgrades & Parts
- Storage Options
- Rental & Leasing Available
- Cycle 5 Master Reseller
- SunSoft Authorized Reseller



To buy or sell,
1-800-253-5764

Ext. 220

Outside U.S. call 1-425-222-7588 Fax 425-222-7388

<http://www.solarsys.com>



Circle No. 411

UNIX Diagramming & Flowcharting

visual thought.

Use Visual Thought for:

- Network diagrams with network component clip art
- Clickable Web diagrams with GIF/JPEG and server/client-side imagemap export
- Framemaker documentation graphics with MIF, EPS) export
- Software design diagrams (Booch, Rumbaugh, Objectory, Fusion, custom notations)
- Flowcharts for ISO 9000, TQM, BPR
- Dataflow diagrams, org charts
- Presentation & business graphics

FREE CD-ROM, FTP & 30-day trial

Additional benefits

- Intelligent, rubberbanding lines
- Macintosh-like ease-of-use... on UNIX!
- Editable, WYSIWYG drag-and-drop palettes
- Hyperlinked, hierarchical documents
- Arbitrary object rotation
- Complete text handling: subscripts, arbitrary fonts, sizes, colors, styles and justification
- Dozens of export formats: GIF89, JPEG, MIF, TIFF, EPSI, XWD, SunRaster, others
- Available on SunOS, Solaris, HP-UX, and soon, Windows 95/NT

VISIT (use code AE13)
<http://www.confluent.com/>
800-780-2838 ext.153
info@confluent.com



Confluent, Inc.
132 Enclave Court
San Francisco, California 94127
415-586-8700 Voice
415-586-8838 Fax

©1998 Confluent, Inc. All rights reserved. Visual Thought, the Visual Thought logo, Confluent and the Confluent logo are trademarks or registered trademarks of Confluent, Inc. All other names are the properties of their respective holders.

See why Visual Thought has thousands of users worldwide! Clients include:

ABB • AT&T • Alcatel BNR • Bellcore • Boeing • CS First Boston • Credit Suisse
E-Systems • EDS • Ericsson • Ford • HP • IONA • Lucent • McDonnell Douglas
Motorola • Nokia • Nomura • Nortel • Persistence • Sprint • Toshiba
U.S. Air Force • U.S. Navy • U.S. West • UBS
Wells Fargo Nikko

Circle No. 447

EXA-"Bitten"??

Extend your EXABYTE Warranty

EXB-820x, 12 hr Repair (6 mo war): \$ 250

"Hot Swap" Repair (6 mo Warr.): \$ 300

EXB-850x, 12 hr Repair (6 mo war): \$ 350

"Hot Swap" Repair (6 mo Warr.): \$ 400

DLT-"Bitten"!!

DLT-2000, DLT 4000, DLT4700: Call

SUN Microsystems W/S

Buy, Sell, Spare Parts, Repairs



West Coast Computer
Exchange, Inc.

11167-A Trade Center Drive
Rancho Cordova, CA 95670

Call: 916-635-9340 or FAX: 916-635-9485

Catalog at Web Site: www.wccx.com

email: sales@wccx.com

Circle No. 453

SAVE \$\$\$

Refurbished Sun Systems & Peripherals

- Ultra & Enterprise Systems
- System Boards
- Drives Memory

Fully Tested Inventory with
same Day Shipping.

Visit our Website @
WWW.COASTALCP.COM



**Coastal Computer
Products, Inc.**

307 Essex Avenue
Gloucester, MA 01930

Tel. 978-282-4488 • Fax: 978-282-4499

Email: wfbaker@tiac.net

Circle No. 348

To Advertise Or For More Information Concerning The

UNIX/NT Recruitment Opportunities Section

Please call Carol Flanagan at (508) 839-4016

NORDISK™
"Your In-Stock Supplier"

**partners
in storage**

NT
HP
Sun
IBM®

**IBM SSA Disk
For NT, HP, Sun and IBM**

- Faster than SCSI
- No single point of failure Available NOW! Call for details and pricing

SCSI Disk

- IBM, Quantum, Seagate, Western Digital
- Capacities to 23GB/Drive

Tape Drives and Libraries

- IBM - 4mm, 8mm and Magstar
- 4mm - HP, Seagate, Tanberg, Tecmar
- 8mm - Exabyte
- 8mm (AIT) - Seagate, Sony
- DLT - Quantum

ADSM-Backup and Data Recovery

- For over 30 platforms
- #1 Rated-The Gartner Group
- Web based administration

Memory

- Kingston (all platforms)
- IBM RS/6000™
- Memory Card Technology

NORDISK™ IBM Business Partner

The IBM logo is a registered trademark and the IBM Partner emblem is a trademark of International Business Machines Corporation and are used together under license. IBM is a registered trademark and RS/6000 is a trademark of International Business Machines Corporation.

Nordisk Systems, Inc.
7900 East Greenway Road, Suite 210
Scottsdale, AZ 85260
Tel: 888-NORDISK or 602-787-9999
Fax: 602-787-9998
e-mail: sales@nordisk.com

Offices: Cleveland • Los Angeles • Orange County • Portland • Scottsdale • Seattle

www.nordisk.com

Circle No. 404

www.nordisk.com **NORDISK™** IBM Business Partner

beat these deals!

NEW Equipment Prices As Low As...

Disk Drives (5 Yr. Warranty)	Memory (Lifetime Warranty)	Tape Drives (2 Yr. Warranty)	CD Rom Drives (2 Yr. Warranty)
• 4.5GB.....\$305	• 64MB.....\$212	Exabyte 7-14GB.....\$1,371	
• 9GB.....\$708	• 126MB.....\$484	Exabyte 20-40GB...\$2,935	Plextor 12X.....\$167
• 18GB.....\$1,451	• 256MB.....\$871	Exabyte DLT4000...\$2,732	Plextor 12/20X.....\$167
• 23GB.....\$1,748	• 1GB.....\$3,339	Exabyte DLT7000...\$6,149	Plextor 32X.....\$190
• IBM SSA.....CALL	• IBM.....CALL	IBM <all types>.....CALL	IBMCALL

sales@nordisk.com

The IBM logo is a registered trademark and the IBM Partner emblem is a trademark of International Business Machines Corporation and are used together under license. IBM is a registered trademark of International Business Machines Corporation.

Circle No. 421

ATLANTIC PERIPHERAL SYSTEMS
"SUPER" SPARC™ VALUES!

MEMORY

64MB SS20.....\$250
64MB Ultra 1.....\$275
128MB Ultra 2.....\$495
256MB Kit Ent. Server.....\$1,095

STORAGE

2.1GB SGT 'cuda int/SCSI.....\$425
4.3GB SGT 'cuda int/SCSI.....\$650
9.0GB SGT 'cuda int/SCSI.....\$950
9.0GB SGT Cheetah int/SCSI \$1,295

SPECIALI

23GB SGT Elite int/SCSI.....\$1,895

WORKSTATIONS

SS Ultra 1, 2.....CALL FOR QUOTE!
NOW! VISIT OUR WEBSITE!
WWW.ATLANTICPS.COM
703-471-0047 FAX 703-471-6621
E-MAIL:atlanticps@mindspring.com

Circle No. 345

Any SUN or HP MONITOR,
will now work with your PC!

The Gemini P1™ is the ONLY graphics card that will drive any fixed frequency color monitor from your PC.

gemini P1 VIDEO CARD

www.si87.com
800-547-2349

Software Integrators Inc. 51 Evergreen Dr. Suite A Bozeman, MT 59715

Circle No. 333

LOWEST PRICES ON MEMORY

- Reliable products backed by free tech. support
- Wide product range for all your systems
- Simple ordering process with fast delivery
- "No Strings Attached" lifetime warranty

Save time, money dealing directly with Manufacturer

Clearpoint Enterprises, Inc.

825 Birch St.
Milford, MA 01757
(800) 253-2778
(508) 473-6111
Fax (508) 473-0112

www.clearpoint.com/~memory
Email: memory@clearpoint.com

"If My Memory Serves Me Right...It Must Be Clearpoint"

Circle No. 346

HP9000 • SPARC • ULTRA • AS/400 • RS/6000 • KSERIES

Get Connected!

KINGSTON

SEAGATE

EXABYTE

QUANTUM

www.wdpi.com

DISKS

CRT'S

CPU'S

MEMORY

FEATURES

TAPE

PRINTERS

Depot Repair • Buy • Sell • Rent • Parts • Training

IBM • HP • SUN

Since 1987 World Data has been providing reliable ontime deliveries and installations of upgrades, CPU's and Features to Fortune 1000 companies around the world. All hardware is guaranteed to be eligible for manufacturers maintenance. What makes us tick? Trained technicians that test all products prior to shipment. Moreover, our reps have the right answers for you and can respond in an instant.

WORLD
DATA PRODUCTS
1 800 553 0592

E-mail: sales@wdpi.com

Fax: 612-476-1903 • 121 Cheshire Lane • Minnetonka, MN 55305 • USA



Circle No. 340

**WTE/WORLDWIDE
TECHNOLOGY EXCHANGE**

- Sun Microsystems
- HP ■ Silicon Graphics
- Workstations/Servers/Peripherals
- Lowest Prices/Extended Warranties
- Will purchase all excess hardware

Tel. 404-378-0990 Fax: 404-378-0939
888-983-7866

Email: sales@worldtx.com www.worldtx.com

Circle No. 322



**INNOVATIVE
COMPUTER SOLUTIONS**

We Specialize in Refurbished Sun Equipment

- ◆ UltraSPARC 1 & 2's, SuperSPARC 5 & 20 Systems
- ◆ Ultra Enterprise E3000, E4000, E5000, & E6000
- ◆ Including Spare Parts & Component Sales

We Also Offer New

- ◆ UltraSPARC Compatibles & UltraSPARC AX Systems
- ◆ Memory Simms, Disk Drives, Tape Drives & Libraries
- ◆ Mass Storage Solutions including RAID Storage Arrays

*FREE Sun Hardware Telephone Support With Sun Trained Hardware Technicians.
On-site Installation Available.*

Call ICS Sales and Consulting 800-838-2850

Email us at sales@icomsol.com

Major Credit Cards Accepted. Leasing available.

Circle No. 331

ENHANCE 3000

Memory manufacturer for Unix/NT servers and workstations

DEC HP IBM INTEL SGI SUN

ENHANCE 3000, INC.

18730 Oxnard Street, Tarzana, CA 91356, USA Phone: 818-343-3066 Fax: 818-343-1436 <http://www.enhance3000.com>

Circle No. 316

WANTED TO BUY

SUN
Cash Paid Now For

ULTRA 1 & 2 • Enterprise Servers
Sparc 5 • 10 • 20s
SUN Disk & Disk Arrays
Memory & Enterprise options
SAVE on OFF LEASE & REFURBISHED SYSTEMS
Systems available for sale, lease or rent
CALL NOW FOR QUOTE



(800) 456-6233
FAX (714) 632-9248

Circle No. 361

HP IBM SUN SGI

BUY - SELL - RENT
NEW AND USED MEMORY

Workstations Servers Mainframes

CYBERTECH STATION Inc.
A subsidiary of Applied Cellular Technology, Inc.
International Memory Resellers

Phone 215.598.9800
Fax 215.598.9867

2288 Second Street Pike, Suite 5B, Newtown, PA 18940

SATISFACTION
GUARANTEED!

Original And
Third Party Products

IBM RS6000 (IBM/IBM)
SUN Sparc, Ultra Sparc,
Enterprise, HP9000/3000,
Netserver and Compaq

Same day shipping on most items.

Circle No. 385

ASR Solutions

Phone: 719-282-3522 Fax 719-282-3523

UNIX computer systems
sales and purchases

BUY > SELL > TRADE > SPARES >

Sun Microsystems, Servers & Workstations
Used & refurbished parts and subsystems
Systems custom configures.

We offer consistently competitive prices and
pay the highest dollar for used SUN equipment.
Same day shipping services available.

Visit us at:
<http://www.quickpage.com/a/asrsolutions>
or e-mail:
asrsolutions@earthlink.net

ASR Solutions

8914 Chetwood Drive, Colorado Springs, CO 80920 USA

Circle No. 362

SERIOUS SUN
SOLUTIONS

Remanufactured Full Technical
Sun Systems Support

Complete Network Services
Peripherals and
Kingston
Memory Upgrades

1-800-842-4781



5980 Golden Hills Drive, Minneapolis, MN 55416
phone: (612) 512-3200 fax: (612) 512-1072
email: sun@workstations.com

Circle No. 320

LIST RENTAL

Are you looking for these people?

- WWW Site Construction/Maintenance Manager
- WWW Site Systems Administrator
- WWW Site Content Developer
- WWW Site Software Developer
- Chief Networking Manager
- Networking Specialist

Search no more. You'll find the above and
many more Web site management selections
in *WebServer Magazine's* subscriber file.

Contact WebServers's highly qualified subscribers
in your next direct mail or telecommunications
campaign. For details call:

Statistics Inc.
Stacie Nestor
(203) 778-8700

BUY - SELL
SUN

& DATA GENERAL

PARTS - DISKS
MEMORY - SERVERS
WORKSTATIONS

AMES SCIENCES, INC

501 SOUTH STREET, EASTON, MD 21601-3846

410-820-8100

FAX 410-820-8179

Circle No. 313

RECRUITMENT ADVERTISING

Reach over

93,000

IT professionals with a recruitment ad in the UNIX/NT
career opportunities section of *SunExpert* magazine.


For special recruitment rates call:

Carol Flanagan at

1-508-839-4016

COMMAND 128 SUNS WITH ONE KEYBOARD & MONITOR


- ◆ Cascadable up to 128 ports.
- ◆ Crisp & clear video at 1600 x 1200 resolution.
- ◆ Comaptible with ULTRA series.
- ◆ Automatic SCAN and keyboard BROADCAST.
- ◆ Hot-plug keyboard and mouse.
- ◆ Our Sun switch can be powered down without causing host failures.
- ◆ Control from keyboard, front panel, optional remote or optional On Screen Display.
- ◆ "Plug-and-Play" - no software required.



FREE CATALOG

Products compatible with SUN, PC & MAC

- ◆ Video Splitters
- ◆ Server Switches
- ◆ Keyboard, Monitor & Mouse Splitters
- ◆ Interactive Classrooms
- ◆ Multi-User Switches
- ◆ Multi-Platform Switches
- ◆ Video Only Switches
- ◆ Cables



NTI Network Technologies Inc
 1275 Danner Drive • Aurora, OH 44202
 800-742-8324 • 330-562-7070 • FAX 330-562-1999
 E-mail: se@networktechinc.com

CALL 800-RGB-TECH
www.networktechinc.com

Circle No. 366



Are your high quality workstation cables costing you as much as your peripherals?

We offer a full line of high quality workstation cables manufactured to meet or exceed OEM specifications.

- **Video cables** - All our 13W3 cables are built with four (4) 75 ohm impedance matched coaxial and four (4) discreet wire pairs. Foil and braid shields, metal cans with 360° soldering and Ferrite beads with contact directly to the shield for maximum effectiveness in reducing EMI/RFI
- **SCSI Cables** - Built to the ANSI specifications with typical impedance on our bulk SCSI cable of 132 ohms. These assemblies have foil and braid shielding and Ferrite beads for maximum effectiveness in reducing EMI/RFI.
- **Serial and Parallel Cables** ■ **Networking Cables** ■ **Fiber** ■ **Custom Cables**

Item	Industry Price	Our Price
13W3 Male/Male Video Cable - 15 Feet	\$180-\$125	\$90
SCSI III to SCSI II with Ferrites - 6 Feet.....	\$125-\$95	\$68
SCSI III to SCSI I with Ferrites - 6 Feet	\$165-\$95	\$68
SCSI II to SCSI II with Ferrites - 12 Feet	\$165-\$110	\$79
Serial Port Splitter Cable - 2 Feet	\$50-\$32	\$23

Attention Resellers and System Integrators!!

We offer a complete Reseller Program designed to assist you with all of your cabling needs.

- One stop shopping for your cabling needs
- Over 400 cables and lengths in stock
- Highest Quality design and manufacturing
- Same day shipping
- Free Technical Support
- Lifetime Warranty

We will help you turn your cable department into a profit center!!



For Ordering Information Call Toll Free
 1-888-6CABLE6 or Fax 1-888-9CABLE1
 1195 Airport Road, Unit 1, Lakewood, NJ 08701
 Phone: 732-942-0086 • Fax: 732-942-0438
www.e-cables.com • email: sales@e-cables.com



Northeast Distributor:
 1120 Ave. of The Americas, 4th Floor • New York, NY 10036
 Toll Free: 800-485-8777
 Tel: 212-626-6785 • Fax: 212-626-6795
www.c-cables.com • e-mail: sales@c-cables.com

Circle No. 422

ADVERTISERS' INDEX

The Ad Index is published as a service to our readers.
The publisher does not assume any liability for errors or omissions.



New England/Upstate New York/ Eastern Canada

CAROL A. FLANAGAN
80 Worcester Street, Suite 9
North Grafton, MA 01536
Phone (508) 839-4016
Fax (508) 839-4226
email: caflanag@aol.com

New York/Mid-Atlantic/ Southeast

JOANNA PARKER
1159 River Road, Suite 2
Teaneck, NJ 07666
Phone: (201) 836-8031
Fax: (201) 836-7958
email: joanna@cpg.com

Mid-West/Southwest/ Central Canada

LINDA LIEBICH
9600 Great Hills Trail, Suite 150 W
Austin, TX 78759
Phone: (512) 502-3035
Fax: (512) 502-9988
email: lindal@concentric.net

Southern California/Arizona/ Nevada/Hawaii/Utah

TARA DUDAS
30 Paseo Alba
San Clemente, CA 92672
Phone: (949) 361-4908
Fax: (949) 361-1564
email: tara@cpg.com

Northern California/Oregon/ Washington/Western Canada

VICKIE MIYAOKA
1935 Mayfield Avenue
San Jose, CA 95130
Phone: (408) 374-9925
Fax: (408) 374-9926
email: vmiyaoka@cpg.com



The Server/Workstation Marketplace Classifieds/Recruitment Ads

CAROL FLANAGAN Mgr. Telemarketing Sales
Phone: (508) 839-4016
MARY ERICSON
Phone (508) 839-0720
80 Worcester Street, Suite 9
North Grafton, MA 01536
Fax: (508) 839-4226
email: caflanag@aol.com

Reader Inquiry Number	page	Reader Inquiry Number	page
439 ...Advantec	96	327 ...Interstar Technologies	98
313 ...Ames Sciences	102	1Kingston	C2
7APC	17	2Lightwave	1
442 ...Apcon	93	372 ...MCA	89
376 ...Ascent Solutions	91	330 ...Michaura Systems	89
362 ...ASR Solutions	102	391 ...MSI	98
19Aspen Consulting	41	343 ...National Data	97
345 ...Atlantic Peripherals	100	374 ...Network Instruments	91
4ATL Products	7	366 ...Network Technologies	103
17Aurora	34	404 ...Nordisk	100
24Aurora	46	421 ...Nordisk	100
363 ...Boffin Ltd.	95	21One Source	42
32Box Hill Systems	BC	353 ...Open Systems Express	89
25Central Data	49	325 ...Frank Orlando Jr.	94
5Chip Coolers	11	13Overland Data	27
346 ...Clearpoint Enterprises	100	318 ...Peripherals Solutions	92
348 ...Coastal Computer Products	99	20Personal Productivity Tools	42
420 ...Computer Connection of CNY	90	18Rave	36
28Concorde Group	59	26Rave	51
447 ...Confluent	99	454 ...Richardson Electronics	95
385 ...Cybertech Station	102	314 ...Security Computer Sales	91
11Data Focus	23	401 ...Security Computer Sales	93
324 ...Datalease	90	438 ...Security Computer Sales	95
455 ...Datalease	96	36SharkRack	47
361 ...Datalease	102	31Siggraph	C3
33DCI	85	333 ...Software Intergrators	100
352 ...Digicad Corp.	96	411 ...Solar Systems	99
428 ...ELI Systems	93	10Summer Internet World	22
34EMASS	9	35Sun Microsystems	63
14EMC	29	27SunSoft	58
316 ...Enhance 3000	101	30SunSoft	84
422 ...Enterprise Cable	103	15Syntax	31
388 ...Evolving Solutions	97	8Tatung Science & Technology	19
.....Exabyte	13	23Transcend	45
431 ...Express Point	95	408 ...Trident Systems	94
378 ...Facet	92	375 ...Ultraview	91
413 ...GSH Systems	90	3Western Scientific	5
409 ...Gulfcoast Workstation	94	320 ...Workstations International	102
12Hitachi	25	16Workstation Solution	33
9Hummingbird	21	340 ...World Data Products	101
317 ...ICCI	90	322 ...Worldwide Technology	101
22Imperial Technology, Inc.	43	370 ...Worldwide Trade Corp	97
29InLine	61	453 ...West Coast Computer Exchange	99
331 ...Innovative Computer Solutions	101	6,40,41,42,43 ZZZZX	15

The World's Premier Computer Graphics Conference and Exhibition Learn about the latest research, art, animation, and emerging technologies in computer graphics and interactive techniques. Enrich your skills with courses and panels. Meet researchers, educators, artists, and practitioners from around the world. Get up-close and hands-on with the newest hardware and software tools from over 350 companies. For complete information and online registration: www.siggraph.org/s98/

discover new worlds



SIGGRAPH 98

celebrating 25 years of discovery


Conference 19-24 July 1998 **Exhibition** 21-23 July 1998 **Orlando, Florida USA** **Contact** +1.312.321.6830 siggraph98@siggraph.org

Titanic TD 27 ©1997 20th Century Fox/Lightstorm Entertainment / The Shape of Space ©1995 The Geometry Center, University of Minnesota / Harvest ©1997 Xaos, Inc. / Project ScienceSpace ©1997 R. Bowen Loftin, University of Houston / Time For Love ©1995 Carlos Saldanha, School of Visual Arts / Rankai ©1997 Nobuo Takahashi / Butterflies in the Rain ©1997 Mark Fontana, Kirk Bowers, and Steve May, The Ohio State University / Virtual Explorer ©1997 Senses Bureau, University of California, San Diego / Multi Mega Book in the Cave ©1997 F.A.B.R.I.CATORS/Italy / Woman Getting What She Wants ©Levi Strauss & Co., Image Courtesy Pixar Animation Studios



Sponsored by ACM SIGGRAPH

Don't be shy because everyone says
you're **smart.**



We're proud of it.

BIGGER, FASTER, SMARTER RAID STORAGE SYSTEMS

Leading-edge technology

Visit www.boxhill.com/raid
for complete specifications on
Box Hill's RAID storage systems and
receive a FREE Box Hill T-shirt.



Box Hill
BIGGER. FASTER. SMARTER.
STORAGE & BACKUP

Circle No. 32

SEE US AT PC EXPO Jacob Javits Convention Center June 16-18, Booth 1540

161 AVENUE OF THE AMERICAS, NEW YORK, NY 10013 PHONE: 212.989.4455 / 800.727.3863 FAX: 212.989.6817 EMAIL: info@boxhill.com WEB: www.boxhill.com