Formerly *SunExpert* Magazine

Server/Workstation
EXPERT

Solutions
for hands-on
IT managers

SOFTWARE
DEVELOPMENT

IN AN OBJECT
WORLD

The Time of SANs

Review: Applixware

YOU GET REMOTE SERVER ACCESS.

YOU GET AN AUDIT TRAIL.

YOU GET TO ATTEND
YOUR KID'S SOCCER MATCH.

Lightwave Communications console management solutions for Unix, PC's andSerial devices fit your needs with a complete line of video,keyboard/mouse, serial (TTY) switches and fiber optic extender kits.

Lightwave's ConsoleServer 3200, ServerSwitch, PC ServerSwitch Plus and SystemConsoleSwitch reclaim valuable server room space and eliminate unneeded peripherals. Remote access puts server control at your fingertips and our audit trail, up to 64K per port, help reduce downtime and diagnose problems quickly. In short, our systems help you reclaim your space, reclaim your time, and most of all, reclaim your life!

Lightwave Communications, Inc.

261 Pepe's Farm Road, Milford, CT 06460
203 878-9838 • 800-871-9838 • Fax: 203-874-0157
Europe: +49 893 063 810
Asia Pacific: +61 396 527 515

Email:sales@lightwavecom.com
World Wide Web: www.lightwavecom.com
Email: office@lightwave.de
Email: sales@lightwavecom.com.au

Circle No. 2

Server/Workstation EXPERT

The leading publication for server-based computing solutions

Contents

Columns

PAGE 20

PAGE 36

PAGE 44

COMING NEXT ISSUE
UNIX Training

BONUS DISTRIBUTION
For Advertisers of the September Issue:
Fibre Channel 1999,
San Jose, CA

Departments

EDITORIAL

dpryor@cpg.com

Look Back to Look Ahead

The halcyon days of multiyear systems projects developed at a measured pace are all but gone. IT programmers and software engineers have to deliver solutions in weeks or, at most, months. As a result, modular, component-based development has become de rigeur. Perhaps the single most important force behind this movement has been e-business.

Because object-oriented programming can make it relatively easy to develop and roll out projects incrementally when time demands rapid application deployment, it has been a godsend to developers. For example, applications can be developed for a business-to-business extranet, then be integrated with enterprise management, then with customer relations, then be brought together in a complete retail system. End-to-end projects would have taken years before the webbed world we now live in. Also, another benefit of all this site-to-site chattering among applications doing e-business has been the accelerated pace of standards adoption. How long would it have taken without the spur of global competition to settle on ActiveX objects, Enterprise JavaBeans and COM/CORBA? For an update on these issues, as well as some ideas about what's next, take a look at Karen Watterson's cover story "Software Development in an Object World," Page 50.

Also this month, *SW Expert* revisits SANland in "The Time of SANs," Page 58. With this feature, Ron Levine takes you into the not too distant future. He discusses the evolution of the storage area network with an eye toward providing the perspective implementers will need through the next two or three years. Of course, he begins with a review of the storage disarray that begat the SAN–businesses relying on a mix of UNIX and Windows NT, performance bottlenecks, scalability woes and so on. But those times are a changing. As those times change, we would like to hear from you regarding your storage plans. If you are planning a SAN, let us know how you're doing it. Nothing helps to anchor our coverage more than first-hand reports from the field. Drop us a line at editors@cpg.com.

Doug Pryor

Exam Time for Java Programmers

Your programmers may be certified for Java development, but are they certified for enterprise-level development? Probably not, because until now the main certifier of Java programmers, Sun Microsystems Inc., didn't have a certification program for enterprise Java development. Previously, Sun Educational Services offered just two levels of Java certification: Certified Programmer and Certified Developer. However, Sun has recently begun addressing the need to test developers not only on their ability to code Java programs and applets, but on their skills using Java services and tools to build distributed, enterprise applications as well.

To that end, in February, Sun unveiled the Java Technology Architecture Planning and Design examination aimed at testing abilities in areas such as database integration, security and high-performance design. More important, in May, Sun and several other vendors announced a collaborative effort to create a joint Certified Enterprise Developer certification. The partners (Sun, IBM Corp., Oracle Corp. and Novell Inc.) plan to have the $125 to $150 certification examinations available by the end of the year, says Steve Holbrook, Java strategist for Novell, Provo, UT.

According to Holbrook, a unified certification will encourage more developers to get qualified in enterprise development. "We wanted something to act as a catalyst for that and we felt it would be more useful for the Java programmers out there if we [vendors] have as much in common as possible," Holbrook says. He predicts there will soon be a lack of qualified enterprise-level Java programmers in the industry. "We're expecting a dearth of programmers as Java spreads, and we won't be able to train people

quickly enough," he says. In fact, Bill Richardson, vice president and general manager for Sun Educational Services, says that according to Sun estimates, the demand for qualified enterprise Java developers already outstrips supply by about 10%.

The joint exam will focus on the APIs and services that form part of

Sun's new Java 2 Platform, Enterprise Edition, as well as on the use of Java development tools. Each of the participating vendors will craft a version of the exam based on its own development products, and applicants can choose which variation of the test they want to take. Because all of the vendors have agreed to recognize each other's version of the test, applicants who test on one vendor's products can still qualify to take an exam for a higher level certification on another vendor's products. So, for example, a programmer who takes the Enterprise Developer exam using Oracle's JDeveloper can go on to take a higher level Java test from IBM without hav-

ing to retake the original Enterprise Developer exam using IBM products.

"We are looking for a way to coordinate [the certification] so someone who comes from one tool base and moves to another tool base doesn't have to start over," Richardson says. "Some of the test is specific to the application tool sets, while other parts deal with the theory or technology basis."

Each company will retain the license for its own test, but other vendors will have the right to offer it along with their exams. It's likely other vendors will join the collaboration in the near future, Richardson says. "There are some in discussion right now. We're pleased to have more participants; the more the merrier."—*sjh*

I/O Specs Race for Pole Position

As CPUs get faster, storage needs expand and peripherals proliferate, the need for faster I/O technologies grows as well. Industry manufacturers are racing to develop new, speedier I/O technologies, but this may result in

overlapping or competing standards that muddy the waters for OEMs and customers.

Currently, three new I/O specifications claim to be the next trend in I/O technology; all promising high-speed performance, and all backed by different computer or networking vendors.

One such specification, Future I/O, is supported by the Future I/O Alliance (http://www.futureio.org), a consortium of more than 100 vendors. Formed in February, the alliance aims to develop a specification for a single interconnect that can be used for both CPU-to-CPU communications in parallel application clusters and high-bandwidth server technologies such as SCSI, Fibre Channel and Gigabit Ethernet. The alliance, which includes Adaptec Inc., Compaq Computer Corp., Hewlett-Packard Co. and IBM Corp., plans to have a final specification out by the end of the year, with products ready for shipping by early 2001.

Future I/O is intended to complement, then eventually supplant, a second (also new) specification, Peripheral Component Interconnect (PCI-X). PCI-X is designed to increase the shared PCI bus to speeds of 1 GB/s. The group working on PCI-X, the PCI Special Interest Group (http://www.pcisig.com), also includes Future I/O supporters such as Compaq, HP and IBM. It plans to have products available later this year or early next year.

"Just as in the early 1990s, ISA [the Industry Standard Architecture bus for PCs] was supplemented and replaced

Java for the Masses

JavaServer Pages (JSP), Sun Microsystems Inc.'s latest Java creation unveiled in June, allows designers to create dynamic, Java-based Web pages without having to know any Java programming. JSP-based pages make it possible to embed Java scripts directly into HTML code. This enables designers to work with the HTML in a Web page, while leaving the business logic that's behind the page to programmers.

"It's a way of inserting real Java code into your HTML," says Bill Roth, product line manager for the Java 2 Platform, Enterprise Edition at Sun, Palo Alto, CA.

JSP-based pages, which are compiled into servlets for execution, require a Web server that supports both Sun's Java Servlet API and JSP 1.0 specification in order to run. Already, several Web and application server vendors, including Bluestone Software Inc., GemStone Systems Inc., Netscape Communications Corp. and Oracle Corp., either are, or soon will be, including support for JSP in their products. Sun has also announced it will license the technology to the open-source Apache Web server developer community. All told, approximately 20 companies collaborated with Sun on the JSP 1.0 specification, Roth says.

JSP competes with Microsoft Corp.'s Active Server Pages (ASP) technology, which dynamically creates Web pages containing either Visual Basic or JScript (Microsoft's version of the JavaScript scripting language) code. Roth says JSP is better than ASP because JSP-based pages are compiled into servlets only once, whereas ASP-based pages must be reinterpreted each time they're run. "Most people building Web applications who want to use Java will want to use JavaServer Pages," says Bob Bickel, senior vice president of products at Bluestone Software, Mount Laurel, NJ. "It's really the simplest way to create Java applications."

However, Bickel says JSP isn't ideal for very large applications. "If I embed all my JDBC [Java database connectivity]

JSP is better than ASP because JSP-based pages are compiled into servlets only once, whereas ASP-based pages must be reinterpreted each time they're run.

code in the JavaServer Pages, and I have an application that has 50 different screens and my database logic changes, then I have to go back into those 50 different applications to change the piece of code in each one. But that's something a fair amount of people are willing to accept, especially for noncritical types of applications or smaller types of applications," Bickel says, adding that JSP is still a good solution for most Web applications.

Flashline.com Inc., a Cleveland, OH-based company that sells software components online, converted its Web pages to JSP several months ago, and now relies on JSP with JavaBeans to automate electronic commerce activities such as shopping cart and real-time credit card processing functions. Charles Stack, chief executive officer of Flashline.com, says the transition has improved both the site's performance and ease of maintenance.

"The huge advantage [of JSP] is it completely separates the code from the presentation, and you get a lot of advantages from that," Stack says. "Your designers can do their HTML design without having to do much with the code, and your programmers don't have to do any HTML whatsoever, which makes everybody happy. And it's faster than CGI scripts, where each page call is a separate process. [JSP] is simply another thread and, as a result, it's significantly more efficient."

Steve Wilcox, chief architect with Avitek Inc., a custom Java development firm based in Boulder, CO, says JSP will relieve programmers of the need to develop their own solutions for generating Web pages. "Prior to JSP, people were using their own homegrown ways of generating HTML from Java. JSP is useful in that it will help standardize that activity," Wilcox says. "We had our own package for HTML page generation, and a lot of people had done the same sort of thing. We no longer have to write our own; we can use something written by other vendors. Sun is filling some holes that existed on the server side."—*sjh*

by PCI, I expect the same thing to happen again," says Martin Whittaker, manager and senior architect in systems and technology for HP, Palo Alto, CA. "PCI-X products will be coming out this year to give a boost to PCI, which will serve as a high-bandwidth bus for about a year and then coexist for quite a while with Future I/O."

Unlike PCI-X, Future I/O will use switched fabric, serial point-to-point connections that offer speeds of 2 GB/s. Switched fabric is not a brand-new technology, says James Gruener, analyst for market research firm The Aberdeen Group, Boston, MA. "It's proven to have worked in the past if you look at mainframe environments, so it's quite likely this will translate well as a workable architecture," Gruener says.

Yet a third I/O technology, the Next Generation I/O (NGIO) interconnect, proposed by the NGIO Forum (http://www.ngioforum. org), will also use a switched fabric serial architecture. The NGIO Forum plans to have products out by mid-2000 and also has a number of

formidable vendors behind it–namely, Dell Computer Corp., Intel Corp. and Sun Microsystems Inc.

"Both [Future I/O and NGIO] have attributes that make them advantageous. Clearly Intel, as well as Sun and Dell, have engineers who under-

stand the architecture quite well. But on the other side, you've got more market size," Aberdeen Group's Gruener says. "The Future I/O folks hold over 60% of the market."

However, Brad Day, analyst with Giga Information Group, Cambridge,

MA, says Intel's presence in the NGIO Forum could serve as a catalyst to drive many OEMs to lend their support to the Future I/O initiative.

"The idea of Intel moving out of microprocessors and into I/O and then into god knows what, is going to bother OEMs enough that they're going to probably have a foot in multiple camps," Day says. "So what may end up happening is that I/O will become just like other standards areas, where some vendors will foot the bill to be able to say they support them all. That's a very expensive proposition."

Brian Smith, chief executive officer of Crossroads Systems Inc., an Austin, TX-based vendor of Fibre Channel storage routers and a member of both the NGIO Forum and the Future I/O Alliance, says that a common specification might be a possibility down the road. "I see Crossroads having a separate set of products in the NGIO and Future I/O spaces, but I also hope there will be a blending of the two [specifications] over time."–*sjh*

App Server Road Map

When Sun Microsystems Inc. agreed to a licensing deal with America Online Inc. of the products AOL acquired from Netscape Communications Corp. last year, industry observers wondered how the two companies would handle product redundancy with their application servers. Sun had recently purchased NetDynamics Inc., maker of the NetDynamics application server, and AOL had released Netscape Application Server (NAS). Now the details are out.

The Sun-Netscape Alliance is taking a two-step approach, with the eventual goal being the delivery of a single application server sometime next year. The first step has already been taken with the release of updated versions of the two existing servers. Both products, NetDynamics 5.01 and NAS 4.0, received a common programming model through the Java 2 Software Developer Kit (SDK) and support for Enterprise JavaBeans (EJB), JavaServer Pages (JSP) and the Java Servlet API. These updated application servers will be followed by further enhancements early next year that will provide shared components in the areas of management and enterprise integration, the alliance says.

The alliance is basically taking what it deems to be the most important features of each application server and integrating them with the eventual single server, says Martin Marshall, director and research analyst at Zona Research Inc., Redwood City, CA. "Their objective is to take the best of both worlds between NetDynamics and NAS," Marshall says. "This is a necessary step for them."

The second step will be the introduction of a new application server, which the alliance says will employ both the development tools of NetDynamics and the transaction processing strength of NAS. This product is scheduled for release in the second half of 2000.

Reaction to the application server road map has been generally favorable, but industry watchers predict the release date of the single server could scare away potential customers. "I think this is going to slow down sales," says Sally Cusack, analyst with International Data Corp., a Framingham, MA-based market research company. "Current customers seem happy, but the end user that is evaluating the technology will wait to see what happens."–*ptc*

Sun Consolidates Software Divisions

Sun Microsystems Inc.'s software employees moved into closer quarters in July, when Sun consolidated its three software divisions–Solaris, Java and Consumer/Embedded–into a single division. The new Software Products and Platforms division will be headed by Alan Baratz, former president of the Java software division.

The consolidation was a practical step taken in response to the recent reassignment of more than 1,000 Sun software staffers to the Sun-Netscape Alliance, established earlier this year, according to John Loiacono, vice president of brand marketing at Sun, Palo Alto, CA.

"In order to increase our software presence, it was only prudent to bring [the remaining employees] together and put even more emphasis on our software products, technologies and solutions," Loiacono says. "To enhance the focus we have placed on software, we decided to be more efficient by aligning all software efforts under a more unified direction and one leading person, resulting in a more unified reporting structure."

The move is a sound one, says John McCormick, strategic analyst for IT software at Frost & Sullivan, a consulting firm based in Mountain View, CA. "Sun had a lot of efforts going on that were redundant and, by consolidating the different businesses, they're able to save some of those costs," McCormick says. It may also give the software side of the business more clout in a largely

The move may give the software side of the business more clout in a largely hardware-driven organization.

hardware-driven organization, he says. "There's strength in numbers."

But McCormick also wonders whether the consolidation might send the wrong signal to Java developers and Sun partners who are participating in work on the Java platform. "Sun's Java platform is a different fish than the Solaris and embedded platforms," he says. "It's kind of an industry standard, and one of the worries people have had about it has been whether Sun intends to be the dictator of that standard," he says. "Consolidating the Java platform stuff with the other software platforms and products puts a tie in there that some people might not like to see."

But Bob Bickel, senior vice president of products at Bluestone Software Inc., Mount Laurel, NJ, doesn't believe that the restructuring will have much of an effect on Sun's current relationship with Java software developers. "There has always been concern, on the part of people like us, about the fact that [Alan Baratz] has responsibility for both the Java standard and for creating profit out of [Sun's Java] software," says Bickel. "But this brings zero changes from our perspective, because both the commercial application side and Java standards group already reported to Alan Baratz [before the reorganization]. This just says that Alan's empire has expanded to include Solaris."

Jean Bozman, analyst at International Data Corp., Mountain View, CA, says the consolidation may even benefit Sun's partners. "It makes it easier for business partners to deal with Sun; they won't have to make as many

trips to talk to Sun," Bozman says.

In the long-run, the restructuring is a good administrative maneuver, Bozman says. "It's an organizational statement, more than anything else, aimed at making software a greater component of their revenues. Sun is simply saying, 'We want to have all the software under one umbrella to manage it better.' At the moment, the vast majority of their revenues come from hardware and services. But it could be, that by doing this, they would become more prominent as a software company," Bozman says.

In another restructuring effort, Sun has created a new Network Service Provider division, which is chartered with developing and marketing Sun products to telecommunications companies, cable operators and network equipment suppliers who develop products for service providers. The division will be headed by John McFarlane, former president of the Solaris software division. Specifically, the division will work on things such as developing integrated software and hardware products, real-time network platforms, network management platforms, extensions to the Java 2 Platform for service providers and fault-tolerant hardware platforms for the telecommunications industry.

Realistically, it will initially act more as a liaison between other Sun divisions, marketing those products to Sun's service provider partners and customers, says Frost & Sullivan's McCormick. "I don't think they're going to go off and immediately start making a bunch of new products. More likely, they're taking Sun's existing stuff and putting a telecom spin on it," he says.

But the move does underscore Sun's intention to be a key player in the emerging Internet service provider (ISP) and application hosting markets. "It

signals a seriousness in attacking that market," says McCormick. "It's kind of interesting they'd have a whole different division. I guess they just wanted to have a unified focus for the market."–*sjh*

Rack 'em Up

A new breed of specialized servers are hitting the market in greater numbers. Instead of powerful, high-performance machines, Sun Microsystems Inc., IBM Corp. and Compaq Computer Corp. are placing an emphasis on functionality and price by offering single-processor systems suitable for specific tasks or markets.

Both Sun and IBM are offering thin servers targeted at the Internet service provider (ISP) market. Sun recently announced the Netra t1–which is available now and has a height of 1U, or 1.75 inches, and a starting price of $5,000–while IBM is preparing for an autumn launch of a similar machine, code-named Pizzazz. IBM's new server

with a height of 2U, or 3.5 inches, will join Big Blue's RS/6000 server family and will be priced lower than Sun's Netra t1, the company says. Both are specifically designed to conserve space.

"One of the chief requirements that has come back from service providers is the need for rack-mounted servers," says Al Rosen, program director of ISP business development for the IBM Server Group. "By stacking them in a rack, and as thin as you can get them, they conserve on floor space, and floor space cost tons of money."

The Netra t1 is available in two versions: a 360-MHz UltraSPARC-IIi processor with 1-MB external cache and 64 MB of memory, and a 440-MHz UltraSPARC-IIi processor with 2-MB external cache and 256 MB of memory. Both boxes run Solaris 2.6/7.

Pizzazz, on the other hand, will be a RISC-based server that supports AIX

and Linux. While IBM is not ready to announce specific features, the company promises it will offer more than Sun's Netra t1, and for less money. "I look at the Netra t1 and it's not a very good performer," says Rosen. "It's not a very interesting price point and I know Pizzazz will significantly outperform it."

At a height of just 1.75 inches, Sun's new thin-client Netra t1 server for the ISP market is designed to conserve space.

Another IBM rack-mounted system in the works is a storage device code-named Oyster. The new storage system, which the company hopes to make available by the end of the year, will provide 720 GB of disk space within a rack-mount enclosure. "You will see us come out with a series of products that fit

CORBA and EJB Integration

Ireland-based Iona Technologies Inc. is attempting to unite Enterprise JavaBeans (EJB) and Common Object Request Broker Architecture (CORBA) with the release of its new development tool OrbixHome.

EJB, a component model created by Sun Microsystems Inc., allows developers to build distributed Java applications with reusable components. CORBA, a specification defined by the Object Management Group (OMG), provides a system for distributing component models among multiple programming languages and platforms. Iona says the EJB component model and the CORBA architecture are a powerful combination.

"EJB gives you ease of use," says John McGuire, senior product manager for Java middleware products at Iona in Dublin. "But when it comes to actually rolling out systems that handle millions of transactions a day, CORBA is a perfect way of doing that."

To bring together EJB's ease of use and CORBA's delivery method, Iona took Sun's EJB 1.1 specification, code-named Moscone, and put a CORBA engine beneath it. Graphical wizards were also added, which allow developers to build custom components or manage ones prebuilt by independent software vendors (ISVs). Iona says these graphical tools can perform all development and management functions, including extending EJB components to CORBA and adding advanced services such as transactions. The CORBA engine provides the messaging and middleware infrastructure.

One key strength of OrbixHome, according to McGuire, is the fact that Iona has based it on the Sun EJB 1.1 specification. "There is an emphasis in the new Moscone specification on assembling beans and deployment support," he says. "So what we are providing with OrbixHome is not just the development capability, but also the infrastructure around assembly and deployment. It comes back to the EJB 1.1 specification with the very tight description of the various roles, what these rules do and providing the tools for that."

Iona, which is known for providing CORBA middleware, paved the way for its EJB tool with the February purchase of EJBHome Ltd., a U.K.-based EJB software firm. OrbixHome is the result of the merger, combining the two companies' technologies.

Iona plans to release OrbixHome by the end of September and will announce pricing at that time. It will be available on Solaris, HP-UX, AIX, True64 UNIX and Windows 95/98/NT. Furthermore, it will support databases from Oracle Corp. and Microsoft Corp. and feature utilities that allow rapid generation of entity beans, a type of EJB that provides a function to map data sources to Java classes.

In addition, OrbixHome will support Sun's Java 2 Platform and is designed to integrate with Iona's Orbix product line. As a result, EJB could be used to drive a remote Microsoft Excel spreadsheet using OrbixCOMet, or to provide a secure Internet firewall, including access control, for EJB collection using OrbixWonderwall.–*ptc*

inside racks," Rosen says.

Compaq's offering is also a thin server designed for the enterprise and ISPs with limited space. In mid-July, Compaq introduced its TaskSmart C-Series towers. With a height of 5 inches, the new Compaq servers are larger than the Sun and IBM machines, but still fit into a traditional 19- by 24-inch telco rack.

Designed specifically for Web caching, and the first of an emerging product category called server appliances, Compaq releases its TaskSmart C-Series server.

TaskSmart C-Series is the first of several new servers to be released that fall into an emerging product category called server appliances. They are turnkey machines, dedicated to a single application or task. For the TaskSmart C-Series, that task is Web caching. "We feel there are select applications for which customers and resellers don't want to be bothered with putting together the optimal solution," says John Young, director of appliance and communications servers at Compaq's Industry Standard Server Division. "They'd rather rely on our experience and integration work."

Other TaskSmart server appliances, which are scheduled to be introduced over the next 12 months, will focus on file storage, security and email. "Very routine, low-customized applications are required for the application format," Young says.

The TaskSmart C-Series server sits between the client and the Web server, caching frequently accessed Web pages. The server appliance can be configured and managed with a Java-enabled Web browser and supports SNMP, Telnet and URL filtering. It's available in three versions: the C1200R, with two 10/100-Mb/s Ethernet ports, 256-MB ECC SDRAM memory and one 9.1-GB hot-pluggable hard drive; the

C1500R, which comes with three 10/100-Mb/s Ethernet ports, 512-MB ECC SDRAM memory and two 9.1-GB hot-pluggable hard drives with support for disk cloning and mirroring; and the C2000R, which comes with redundant hot-pluggable power supplies, five 10/100-Mb/s Ethernet ports, 1-GB ECC SDRAM memory and six 9.1-GB hot-pluggable hard drives. All three models run caching technology from Novell Inc. and a stripped down version of the NetWare 5 kernel. At this writing, Compaq has not released prices for the TaskSmart C-Series servers.

One of the first users to test the TaskSmart C-Series server was Don Porter, senior network engineer for UtahLink, an online organization that manages Internet traffic for 41 public school districts throughout Utah. Porter is running TaskSmart as a proxy cache to filter content intended for the students. Previously, he was using Sun boxes to do the same function. "I had approximately 12 Solaris boxes scattered across the state. If we had a problem, sometimes we couldn't get someone on-site to have the expertise to resolve the issue," Porter says. "[With TaskSmart] I just provide the end user with the hardware, the CD and floppy, and if they have a problem they just rebuild."*–ptc*

Rival XML Sites Launched

Despite current clear skies, a potential tempest is brewing in the world of eXtensible Markup Language (XML). In late May, two competing Web sites were launched within 24 hours of one another, both with the intention of promoting the use of XML-based business schemas.

XML.org, unveiled on May 25, is an XML portal developed by the Organization for the Advancement of Structured Information Standards (OASIS), a nonprofit consortium backed by industry leaders, including

IBM Corp., Sun Microsystems Inc., Oracle Corp. and Novell Inc. The other, BizTalk.org, an effort led by Microsoft Corp. and used to promote its proprietary XML framework, was launched the day before. The introduction of the BizTalk site has led to speculation that the emerging XML market will soon develop into a power struggle to create the leading Document Type Definitions (DTDs) and business schemas. "It looked like there was a battle royal shaping up between Microsoft and OASIS," says Steve Robins, senior analyst at The Yankee Group, a Boston, MA-based research firm.

XML is a vendor-neutral data exchange language being developed by the World Wide Web Consortium (W3C), which allows for structured document interchange on the Web. Because many industries have their own specific rules, unique business schemas are being designed for XML data exchange. BizTalk is Microsoft's version, while other industry organizations and vendors are attempting to define their own schemas. One fear is that a company with the clout of Microsoft can establish the standard that best fits its proprietary needs and not that of the entire industry.

However, those concerns were placated–at least for the moment–when Microsoft informed OASIS on May 28 that it intends to join the group. "We have long sought Microsoft to be a member of OASIS and we did speak to them about our XML.org effort several weeks before they had launched their [BizTalk.org] site," says Laura Walker, executive director of OASIS. "But even if Microsoft had not joined OASIS, the two efforts were really complementary and not competitive."

Walker says the two sites provide a similar service. She hopes XML.org will become the portal leading to all things XML. That includes providing the different business schemas and DTDs that are developed in the industry, as well as ones developed for the BizTalk Framework. "Companies exist to make money and it is through product differentiation that they are able to do so. What we focus on at OASIS is that even though there will always be

When it absolutely positively has to stay up all the time.

RESILIENT

Non-Stop Computing

360 MHz UltraSPARC

Firewall, DNS, Web...

Available in Rack Mount

Introducing the Resilient® Ultra fault tolerant server.
It delivers when any downtime or data loss is too much.

We created the Resilient® Ultra fault tolerant server specifically for your mission critical network computing applications: firewalls, web, DNS, NIS+, license server, telecom, databases — any job where even a little downtime is too much.

Unlike software-based "high availability" schemes, our triple redundant architecture keeps applications running non-stop even if a logic board fails. Which means you won't experience unrecoverable data loss, service interruptions, high administration costs, reconfiguration delays — or constant worrying about your critical applications.

Resilience fault tolerant servers are based on industry-standard Sun UltraSPARC technology and the Solaris OS, so they integrate seamlessly with your existing network and applications. Their hot-pluggable modular construction offers exceptional ease of maintenance — you can even replace any failed component in under 60 seconds.

So now you can enjoy non-stop availability and the performance of an UltraSPARC server — at a surprisingly low cost. For an informative white paper with more information about the new Resilient Ultra, simply call us at (800) 977-9008 or visit www.resilience.com.

·RESILIENCE
Fault-Tolerant Network Computing™

1755 Embarcadero Rd., Suite 120 · Palo Alto, CA 94303 · (800) 977-9008 · (650) 843-3600 VOICE · (650) 843-3609 FAX · email: sales@resilience.com
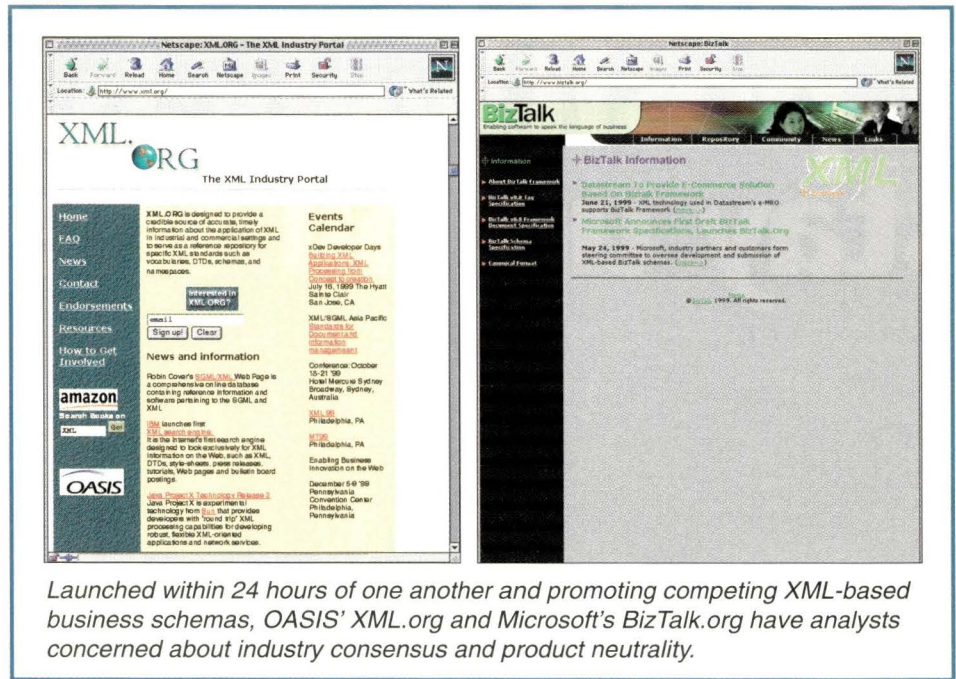
Circle No. 8

this product differentiation, the exchange point between products is neutral. That is what we strive for," Walker says.

Walker's position is a common one among leaders in the XML industry. In fact, many believe there will be numerous types of business schemas and DTDs. "The reality is there is just a need for a variety of different DTDs and schemas out in the world," says Bob Bickel, senior vice president of products at Bluestone Software Inc., a leading vendor of XML products based in Mount Laurel, NJ. "The world is just too complex to try to model in one DTD and schema."

Most people feel it's better to have Microsoft onboard with XML.org with the hope that industry consensus can be reached on the different XML schemas being created. "This is extraordinarily good news for XML," says The Yankee Group's Robins. "XML has been very strong as a way to share data content. What has made it a challenge is different schemas for sharing that data. But instead of there being different organizations propagating these DTDs, it

Launched within 24 hours of one another and promoting competing XML-based business schemas, OASIS' XML.org and Microsoft's BizTalk.org have analysts concerned about industry consensus and product neutrality.

is really going to be one."

Even Microsoft's biggest competitors believe its joining OASIS is a positive step for XML. "We're really happy to see that Microsoft has signed up," says Nancy Lee, senior product manager for XML at Sun, Palo Alto, CA. "It's definitely something that is good."

While Microsoft did not respond to requests for an interview regarding its

plans to join OASIS, it is safe to say the company does not plan to abandon BizTalk.org and its BizTalk Framework. If that's the case, concern over industry conflicts will undoubtedly continue. "With the presence of [BizTalk.org] there is the possibility Microsoft is promoting a proprietary version of XML," says Lee. "That could be an issue."–ptc

READER FEEDBACK

To help *Server/Workstation Expert* 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.

	Interest Level		
	High	Medium	Low
Features:			
Software Development in an Object World	170	171	172
The Time of SANs	173	174	175
Product Review: Applixware: A Fine Alternative	176	177	178
Columns:			
Ask Mr. Protocol–Mr. P. Brings Home a Movie	179	180	181
UNIX Basics–What Machines are out There?	182	183	184
NTegration–Upgrade Like it's 1999	185	186	187
Work–Software Ptools	188	189	190
Java Class–Across the Great Divide	191	192	193

THE MOST IMPORTANT IMAGES OF THE 20TH CENTURY ARE BACKED-UP ON QUALSTAR TAPE LIBRARIES.

TOM BARRETT

"Somehow Julia Roberts isn't quite so attractive when she's made of big square blocks like that."
– A cable TV viewer

"Email services will be restored as soon as we figure out what went wrong."
– A cable modem provider

"Let's rent a movie. Uh, how much space do we have on our hard drive?"
– A future viewer

Mr. P. Brings Home a Movie

Q: *I read in the paper that movie studios are already starting to look at the notion of distributing movies over the Internet. I couldn't care less about that because my TV works very nicely, thank you, but I'd love to be able to transfer movies over the Internet on my own. Only my movies are sales presentations and training films for my company. How do I do that?*

A: Your question raises so many others that I don't know where to begin.

First of all, the movie studios are indeed studying the idea of delivering movies over the Internet, but only partly because they want to. Pirate sites are already delivering very bad copies of movies over the Internet, mostly made by semi-bright people sneaking digital camcorders into movie theaters. This works about as well as you'd expect. If these were any worse, they would look like videotaped evidence used in court.

The studios are in the best position to deliver high-quality movies via the Internet, but encounter the same problem they face with videotaped movies: piracy. If they took the simplest solution, namely, running something like an FTP server that took credit cards, sales would be decimated by pirates selling cut-rate copies 10 minutes after a movie was released. Somehow, they have to be able to deliver digital copies of movies that cannot themselves be copied, or at least, not usefully copied.

Mr. Protocol thinks he's found an approach that'll do the job. Several companies have announced digital sales schemes, which deliver the digital product wrapped in a cryptographic holder. These wrapped products can be made freely available because as they stand, they're useless. What you buy is a key that decrypts the product, but only under certain conditions. If the product is a movie, the key may only decrypt it for a single viewing, or a limited number of viewings, and only on the user's own machine.

Mr. Protocol's tenuous grasp of physical reality has caused him to ignore one potentially fatal flaw in this approach: Any encryption scheme requires that the item in question exist in decrypted form in the user's own domain, however fleetingly. Preventing a sufficiently resourceful user from capturing the product at that point is a difficult problem. A movie, for example, can be captured on its way to the screen, if not earlier.

And, because you're enamored of your own set, a word about TV. Current televisions use a standard called NTSC, which stands for the National Technical Standards Committee, the collection of dead white guys who thought up the standard back in the 1930s or so. Screen resolution is a relative thing, but it's safe to say that just about all computer displays today have far higher resolution than television screens. That's why products like WebTV produce displays, which look very different from "ordinary" displays on a computer monitor. A regular Netscape page would be com-

pletely unreadable on a TV screen, for example. High-Definition Television (HDTV) will improve that picture considerably (no pun intended) but, although HDTV sets are now available, they are not expected to become widespread until the broadcast stations cut over to the new format several years from now.

And high resolution is a winner. Mr. Protocol is aware of someone who recently ran out and bought a high-end Apple Computer Inc. PowerBook G3 notebook computer with DVD drive, and who has since been spending most of her time watching DVD movies on it. They look and sound great. Mr. Protocol believes it. Anyone who's taken a look at the current crop of computer games, which demand 300-MHz processors and 3D rendering cards, will be aware that the real capabilities of current computer displays should not be judged by the muddy, postage stamp-sized movies generated by the current crop of streaming video applications.

The problem is these things are large. Very, very large. Showing high-resolution, smooth video on a computer today requires that the video be stored locally, and that if it is read from a CD-ROM, the CD-ROM drive must be very high-speed, or the data must be highly compressed (as is the case with DVD). Apple's QuickTime video is an exception to the "muddy postage-stamp" rule, because QuickTime videos are downloaded to local disk in their entirety. The QuickTime application won't start to play a movie until enough has been downloaded so there is time to get the rest downloaded while the first part is playing. Real Networks Inc. RealVideo and Microsoft Corp. Windows Media don't have that luxury. They buffer just enough to stay ahead and use a viciously high compression ratio, which often makes the picture look like a confusing blot of muddy colors if the entire picture is showing things in rapid motion.

What can be done about this?

Mr. Protocol is glad you asked.

More Bandwidth, Better Compression

The two obvious answers are more bandwidth and better compression. Compression is a tricky subject, one which only mathematicians are really happy with. The amount of compression that can be achieved "cold" is limited. Even if you achieve 60% compression, you've cut many, many megabytes down to many megabytes, which doesn't do you a lot of good if you have to cram it all over phone lines. To get really good compression, you have to cheat.

The best kind of cheating involves foreknowledge. If, for example, you already have an original 35mm print of *Notting Hill*, you can encode the whole movie as a single bit, with meaning, "show *Notting Hill*." Despite the extremely high compression ratio of billions to one, the result is rendered in gorgeous detail, enough to remind one forcefully why Julia Roberts commands an eight-figure salary, while permitting one to forget that what happens to Hugh Grant in this movie happens just often enough in real life to keep hope alive.

You'd been thinking about a new watch for a long time with the idea of Rolex playing on your mind. Then you remembered grandpa's old timepiece – still packed away since the day he died.

As you lifted it from the box you could feel its strength. There was more character in its handsome face than in anything you could buy new.

So you fixed the mechanism, replaced the crystal and polished it 'til it shined to perfection. Now it's alive and ticking... a beloved family heirloom... your proudest possession.

That's how we feel about computers.

New isn't always better.

When it comes to delivery, timing is everything.

We do everything in our Repower to meet your delivery needs – within a few days... sometimes overnight. The fact is, we can Repower your system, integrate and custom configure it to your specifications and get it to you in record time...usually faster than anyone else. At prices that will take you back in time.

Check us out at www.radiantresources.com for aggressive pricing, net specials – and the chance to configure and instantly price systems in your own good time.

What a difference a day makes

Recently we got a call from an IS guy in a Fortune 100 Company. He needed an E4000 – *at any price*...within 24 hours...by Friday of Memorial Day weekend! And none of his vendors – neither OEM nor VAR – could deliver. You can imagine how wound up he was by the time he called us, especially since he wasn't even our client.

But we turned on the Repower of Radiant, processed his order and hand-delivered it to him in time to meet his impossible deadline. Made him look like a hero. And he's been thanking us with business ever since.

So start saving your company time and money with Repowered equipment from Radiant.

'Cause like the classic elegance of grandpa's old watch, you'll find the benefits of our quality and service utterly timeless.

Call Toll-Free–800-500-7675

Radiant
Resources, Inc.

Repowered Computers For The Next Generation

1913 Atlantic Avenue, Suite R-3
Manasquan, NJ 08736
973-442-5555
Fax: 973-442-5550
www.radiantresources.com

Major Credit Cards Accepted.

Circle No. 24

One-bit compression is an intellectual relation of the mathematical "proof by assumption," in which the result of the theorem is cleverly concealed in the axioms, allowing the theorem to be proved by inspection. Nevertheless, there are compression schemes which reduce to this in extreme cases. Huffman coding, for example, breaks an input string into a series of substrings, so that repeated instances of a given substring are encoded by a potentially much shorter combination of bits, which signals, "Got another one of these."

In the case of an input string consisting of 8,257 repetitions of the letter "e," for example, the input string would be vastly shortened: "There's a substring of 8,257 e's in a row. Got one of those."

Most current video compression schemes involve breaking up each frame into a number of square tiles. The initial, very rough encoding consists of an average color and brightness for each tile, sort of like the famous picture of Lincoln made up of a series of large square blocks of different colors. It doesn't look a thing like him, but it's still recognizable.

If the algorithm is lucky, a given tile will be part of the sky or the ground, and will have little variation within it. In that case, rendering more detail into the square won't require very many bits. Most of the bits will go into breaking down the squares that have a lot going on: people's faces, edges of overlapping objects and so on.

The trick is to come up with a "predictive algorithm," one that makes assumptions about what's going on, and then proceeds to encode, not the real data, but the differences between the data and its own assumptions. Because the decompression algorithm either knows the assumptions, or can be informed of them with relatively few bits, the compression of frames where those assumptions are valid can reach quite high ratios without significant data loss.

One big assumption that's generally made is that successive frames of video aren't wildly different from one another. In such a case, one encodes the first frame fully, then sends only the differences between that frame and succeeding frames. This breaks down relatively infrequently. Very rapid montage shots are the worst, because they violate the axiom most severely: every frame really is different from every succeeding frame, or close to it. Then there are the handheld camera shots, very common in news footage, where the whole picture shakes around. Some recent compression algorithms are still able to handle that; they have a way of encoding the statement: "Start by shifting the whole picture 30 pixels right and 82 pixels up."

The movies shown on cable TV by the pay services like Cinemax and Showtime are compressed in this way. That's why when there's a momentary glitch in the satellite feed, the picture freezes and parts of it break up into colored square blocks. The data stream has been interrupted in midframe and the hardware can only paint what it's got.

Still, there is a limit to what compression can do. It's hard to say what that limit is because increasingly clever compression algorithms are still being developed, but for now, bandwidth is the only way to avoid the muddy postage-stamp syndrome.

One recognizes right off that dial-up connections are useless for video. The only thing that looks good over dial-up lines using today's compression schemes are talking heads at news conferences (and then only if no flags are actively waving in the background). The most likely short-term solution for a compression algorithm to fix this would require that a) the Egyptian tombs were built by space aliens, and b) one of them doodled such a helpful algorithm on the wall of an as-yet undiscovered tomb. One wonders what the determinative sign of the Middle Egyptian word for "video" would be. One doesn't actually want to think about this sort of thing for too long, though.

Bandwidth it is then.

At the moment, the two near-term solutions to higher bandwidth appear to be cable modems from the cable company and Digital Subscriber Line (DSL) service of one flavor or another from the telephone company.

Cable modems generally offer speeds of 1.5 Mb/s coming in and 300 to 600 Kb/s going out. This is fine for most services, so long as one is not supporting a very hot Web site on one's local machine. However, cable modems so far are the exclusive domain of the cable company that owns the wires. As this is written, a federal judge in Oregon has held that the local cable company must open up to competition, but there are arguments on both sides.

Opening up the cable modem market to competition at this stage, reverses an entrenched public policy that's almost 100 years old. When the telephone was introduced, various independent companies provided local service, just as Internet service providers (ISPs) do now. Unlike the ISP case, if you and the person you wanted to call were on different systems, too bad!

The point is the government wanted to foster the installation of the tremendous infrastructure necessary to meet the goal of "universal service"–a working telephone for everyone, no matter where they lived. Only when that goal was well past being met was local telephone service once more opened up to competition.

Many feel that the dynamics and economics of the Internet are similar. In fact, some claim that if the cable companies are required by new legislation to grant access to their cable system to competing ISPs, installation of new cable modem service will grind to a halt.

The problem is that the cable infrastructure that was installed to move television signals isn't up to moving Internet data. As Mr. Protocol has mentioned before, the standards originally set for the cable industry foresaw the eventual need for two-way communications, but the early cable companies, balking at the expense, rammed through an addendum that allowed for one-way systems, and then built nothing else.

Therefore, all of the cable line amplifiers had to be replaced to allow two-way communications. And remember that bit about TV resolution? The coaxial cable used in TV cable systems is miserably cheap stuff. The braided shield, which is supposed to cover the inner insulator and conductor completely, looks more like a fishnet stocking. All of that stuff had to go, too.

Consequently, what the cable companies originally did on the cheap now has to be reinstalled and rebuilt completely. The only thing that doesn't have to be redone is the negotiation for the cable right-of-way.

With this kind of investment being made, companies are loath to then throw it open to all comers. Consequently, a considerable war is being fought in the political arena to ensure that cable Internet systems remain as monopolies. The problem is the original telephone monopoly was a competent monopoly. Cable modem systems, on the other hand, tend to pull stunts that leave one gasping. Days of no mail service. Days of no DNS service. Things that in a dial-up ISP, would send customers positively barreling over to the competition. Only there isn't any. And who wants to give up 170 Kb/s or so for 5.5 Kb/s, maximum? When it does work, it's pretty sweet stuff. So people suffer, and send endless service emails and make endless service phone calls. Of course, the cable companies all have network status Web pages, which is wonderful if your network connection happens to be working, but as with most other ISPs, the network status pages represent a fiction aimed at investors, not customers. ISPs have found that if investors see much trouble reported on the network status pages, their stock falls, so they've learned to report only scheduled maintenance and smooth sailing on these things.

What's the Alternative?

There is a high-speed alternative. DSL service is provided over phone lines by the phone company. The touted advantage over cable (besides working with the phone company instead of the cable company, which is like working with the Roman Curia instead of the Huns–disadvantages both ways) is that in a cable modem system, everyone on a given cable segment is actually on something like a great big Ethernet: a shared medium. Shared packets, shared bandwidth. DSL lines are dedicated, high-speed lines, just like cable modems, but they're individual, not shared.

The reality is that all of the DSL lines go into a router back at the phone company, and the capacity of that router is based on average aggregate use, not peak aggregate use, so it turns out the bandwidth is shared after all. Plus, DSL lines are generally slower than cable modem lines. Full T1 rates over DSL, which match cable modem rates, are far more expensive than cable modems.

Even allowing for the birthing pains of high-bandwidth home Internet connections, there are other problems to be addressed. For one thing, holding even a compressed movie is going to take a lot of bits. Even with today's relatively inexpensive 9-GB drives, a single movie would fill most of the space on your hard drive. But even at T1 rates, streaming video over the Internet still looks pretty crappy. Plus, anyone who's tried it knows that network congestion can bollix up even a short news

clip in mid-transmission. "Network congestion – buffering" is an all-too-frequent message.

There are solutions being worked on for this. On one side, we have the folks who believe that any system that wants uninterrupted bandwidth reserved for its use should be able to negotiate for it. This is the basis of the Resource Reservation Protocol (RSVP). In this protocol, "end systems" set up a bandwidth reservation for a given connection. An end system is either the system sending the data or the system receiving it. RSVP is initiated by the receiving system. It sends out a request for bandwidth over the connection, and each router along the way reserves bandwidth. Once the request reaches the sending system, the sending system sends an acknowledgment back to the receiver and the contract is complete. If the reservation fails at any point along the way, the receiving system knows that the request cannot be accommodated.

Anyone who's ever played with RSVP knows it can be impressive. A video stream that is impossibly jerky over a normal TCP connection, or using UDP, suddenly gets smooth as glass when an RSVP contract is negotiated. Although it is not widely deployed, RSVP seems to work well for point-to-point video streams…so long as all the critical systems in between understand RSVP.

On the other side, there's a body of folks who don't believe that end systems should have to mess with any of this. They have another solution in mind, called "active networks." An active network is one where the routers, instead of just sending each packet that comes along on its way (or dropping it), pay attention to the nature of the traffic. An active network can tell the difference between email being sent in the background, interactive sessions running in real time and video and audio streams, purely from their behavior and the port numbers involved. Exactly how they do this, and what they do when they've got it all figured out, is a current research topic–but it is a very hot topic indeed. The Defense Advanced Research Projects Agency (DARPA) has a major initiative currently underway on active networks.

Unfortunately, which solution one prefers is pretty much a matter of religion, just like the choice of operating systems. The two solutions can coexist, because they operate pretty much at right angles, but the adherents of each can be militant. And while different operating systems can run in different institutions, in the case of the Internet, it's likely only one solution will prevail. It's too early to say which one that will be. ➠

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@cpg.com.

UNIX Basics

by Peter Collinson, **Hillside Systems**

PAUL SCHULENBURG

What Machines are out There?

I've been engaged in some form of computer networking for about 20 years. My earliest experience, in the late '70s, was connecting two computers together using a serial line. I moved on in the early '80s to write the code that was used in the hosts on our campus-wide network based on the Cambridge Ring. In the late '80s, it was back to serial lines (and X25) for the UUCP network.

The '90s have seen the emergence of the Internet and its killer application, the World Wide Web. As we approach the end of the '90s, the Web seems to be turning into something run by big business for big business, whose glossy Shockwaved sites often contain little of interest. Advertising agencies earn their money by turning nothing into something, but because they start with nothing, the actual content of these sites is zero. Nothing in, nothing out. For me, the essential aspect of the '90s Web has been its ability to give you access to information that was not available

before. I hope the delivery of real information doesn't completely disappear in a wave of commercially created electronic billboards.

A very common experience in these 20 years has been the helpful phone call from the person who tells you that your network is broken in some way: your mail system is down, your Web server isn't responding, or some other cataclysm has befallen you. The call usually occurs at an inconvenient moment, and the stress is compounded because you have been using the network connection for some considerable period and were convinced that all was well with the world.

Now, the call is actually well-intentioned, it's someone trying to be helpful, so you cannot be too rude when you discover that the fault is theirs. I think I have grown into a state where I always presume that they have the defective connection and whatever fault they have isn't my problem. I always try to be polite when I point the fickle finger of fate to the deficiencies in their setup.

After all, humility dictates that sometimes, only sometimes, the presumption of my innocence is misplaced.

I recently had an experience where someone was trying to send me mail and called saying, "Your mail system is broken." This is not a terribly helpful thing to be told. I tried not to get cross. "How is it broken?" "It's bouncing the mail that I am sending." "What does it say in the bounced message?" "Oh, I haven't looked at that." The usual cause of this error is a mistyped address and the bounce is actually from the user's Internet service provider (ISP) and not from the target mail system. I presumed my innocence.

The caller found the bounced message and read the error out to me. Further investigation showed that my mail system was indeed bouncing his mail because his domain was in my list of spammers, which meant that at some point in the recent past, a user of his ISP sent me considerable volumes of unwanted mail. I informed him of this,

Listing 1. dig Lookup

```
$ dig www.hillside.co.uk

; <<>> DiG 2.0 <<>> www.hillside.co.uk
;; ->>HEADER<<- opcode: QUERY , status: NOERROR, id: 6
;; flags: qr aa rd ra ; Ques: 1, Ans: 2, Auth: 4, Addit: 4
;; QUESTIONS:
;;      www.hillside.co.uk, type = A, class = IN

;; ANSWERS:
www.hillside.co.uk.     86400    CNAME    wooded.hillside.co.uk.
wooded.hillside.co.uk.  86400    A        194.205.42.3

;; AUTHORITY RECORDS:
hillside.co.uk. 86400   NS       ns.hillside.co.uk.
hillside.co.uk. 86400   NS       nsa.hillside.co.uk.
hillside.co.uk. 86400   NS       ns0.insnet.net.
hillside.co.uk. 86400   NS       ns1.insnet.net.

;; ADDITIONAL RECORDS:
ns.hillside.co.uk.      86400    A        194.205.42.3
nsa.hillside.co.uk.     86400    A        194.205.42.1
ns0.insnet.net. 110796  A        194.177.160.34
ns1.insnet.net. 110796  A        194.177.170.34

;; Sent 1 pkts, answer found in time: 2 msec
;; FROM: craggy to SERVER: default -- 127.0.0.1
;; WHEN: Wed Jun  2 11:17:55 1999
;; MSG SIZE  sent: 36  rcvd: 232
```

and asked him whether he should be using an ISP that supported spammers. He made no comment.

I relate this tale because if you are going to be helpful, then you should realize that the person you are calling will probably assume that their system is working well. You can be more helpful (or plant the seeds of guilt) by transmitting any error messages you have received. Now, I will agree that error messages from mail systems are not exactly user-friendly and often contain loads of gobbledegook that probably frightens novice users, but if you look at the message, there is usually one sentence that explains the error.

On the whole, I try to diagnose an error before complaining about it. It doesn't take a guru to perform some basic checks on a remote machine or the service supplied by a remote machine. Quite often, these checks show there is something wrong at your end, and you can rectify it, or at least route an error report sensibly to get the problem fixed.

UNIX Basics

Verifying an Address

The most common fault is undoubtedly the use of an address that either doesn't exist or where the host machine has disappeared for some reason. As you probably know, domain names are mapped into IP addresses by the Domain Name System (DNS–the "S" can stand for "Service"). I discussed DNS in general terms in an article published in May 1998, "The Domain Name Service," *SunExpert*, Page 30 (http://sw.expert.com/C2/SE.C2.MAY.98.pdf). The DNS distributed database allows me to control my local address space using local files while telling you about it. Basically, let's assume you start an application that wants to look for one of my machines, craggy.hillside.co.uk. The application will ask your local DNS name server and that server will first reach out over the Internet for a server that supports the .uk domain. Then, with data based on the .uk information, it will look for a server that supports .co.uk and then for the address of my name server containing hillside.co.uk. Finally, your code will interrogate my name server to find details about the actual machine craggy.hillside.co.uk. The search can take some time, so the DNS code at your end will retain any results for some period so that a subsequent lookup will return local information, which may actually be out of date.

Also, some applications will give up and say, "Sorry I cannot find that information." This is particularly the case with Web browsers. If someone calls you and says, "Your Web site is down," when you know that it isn't, there's a good chance their browser has simply given up on the chore of looking up your address.

When you ask them to retype the address into the browser window, magically things will spring into life. In the interim, the needed DNS information has arrived and is present on their machine. At one time, I got several complaints about lack of Web service from my site from users of one ISP whose DNS system was simply overloaded. Asking them to type the address again sorted out the problem.

The DNS is a publicly available database, and there are user-level tools that allow you to interrogate it. Most machines will have the nslookup program, which is sometimes a little mysterious (it's hidden in /usr/sbin on Solaris). For quick lookups, I prefer the dig program that emanates from the University of Southern California. It's not a standard part of Solaris, but precompiled binaries are available. On my Red Hat Linux system, and also on my BSD system, dig is sitting there ready to use.

Listing 1 shows you how to lookup a machine with dig. Please pick another address when you try this for yourself, I do not need the extra traffic. The output is somewhat voluminous, and I don't intend to exhaustively explain it all here. The main reply to the request follows the ANSWERS: line. It shows that www.hillside.co.uk exists as a name in the DNS. This is an alias (a CNAME) to my machine wooded.hillside.co.uk, which has an IP address in an A record. The large numbers that follow the machine names are time-out values for the name, and I'll ignore them here. The remaining information tells you about the name servers that support my domain.

Looking up names using dig works for named machines like www.hillside.co.uk. You may think www implies a service, but this is only a convention, not part of the network. What about mail? The DNS supports a special type of record, the MX record, which tells mail systems where mail is to be sent for that address. MX records are used to advertise the names of a range of machines that are prepared to forward mail to the particular mail address. Again, if I use dig to inspect my mail address, I'll get the full answer as shown in Listing 1, but the ANSWERS: will be

```
$ dig mx hillside.co.uk
  ...
  ...
;; ANSWERS:
hillside.co.uk. 86400 MX 10 craggy.hillside.co.uk.
hillside.co.uk. 86400 MX 20 mxbackup.insnet.net.
  ...
  ...
```

which tells the mail system that it should first send mail to craggy and then to the backup mail system supplied by my ISP: mxbackup.insnet.net. The numbers before the machine names ("10" and "20") are priority values, the lower the better.

The MX records in my DNS record allow you to address mail to my site using my domain name. Actually, some people don't use this mechanism, and should. One option is to associate an IP address with your domain name and then hope that the sender's mail system will behave like sendmail. If sendmail cannot find an MX record for the address, but can find an IP address, it will send the mail to the IP address. The thinking here is that you should always be able to send mail to a specific machine. However, people have perverted this thoughtful mechanism to allow the transmission of mail to domains that happen to have IP addresses. Associating an IP address with a domain name is actually discouraged by the Internet standards, but people still do it.

The final point to make about dig is that it can be used to painlessly find the reverse mapping, translating an IP address into a machine name:

```
$ dig -x 194.205.42.1
  ...
  ...
;; ANSWERS:
1.42.205.194.in-addr.arpa.
              86400 PTR craggy.hillside.co.uk.
  ...
  ...
```

This shows that my domain has reverse mapping entries (PTR records) that map onto the fake domain that is used for reverse lookup. Again, I am finding that some ISPs don't set up the PTR records properly and, increasingly, you cannot find a machine name for a particular IP address. My FTP system refuses to deal with such anonymous IP addresses, and I encourage you to do the same. I support anonymous access, meaning you

---

Listing 2. Using ping

```
$ ping -s craggy.hillside.co.uk
PING craggy.hillside.co.uk: 56 data bytes
64 bytes from craggy.hillside.co.uk (194.205.42.1): icmp_seq=0. time=169. ms
64 bytes from craggy.hillside.co.uk (194.205.42.1): icmp_seq=1. time=165. ms
64 bytes from craggy.hillside.co.uk (194.205.42.1): icmp_seq=2. time=272. ms
64 bytes from craggy.hillside.co.uk (194.205.42.1): icmp_seq=3. time=236. ms
64 bytes from craggy.hillside.co.uk (194.205.42.1): icmp_seq=4. time=331. ms
64 bytes from craggy.hillside.co.uk (194.205.42.1): icmp_seq=5. time=348. ms
^C
```

---

don't have to have an account, therefore access is anonymous. But I don't support your anonymity. I don't see why I should not know who has visited my system and pulled files.

The ping Command

Well, as you can see, `dig` allows you to discover how names map onto a particular IP address, and as a side effect, allows you to check that you're using a legal address. What if the address you are using is correct, but things are still not functioning? What next?

Well, is the machine that you need to access up and running, and can you get to it? The `ping` command can help. The command sends a message to the remote host and looks for a reply. If the remote machine replies, then it's alive. So on a Solaris or SunOS system, you'll see something like the following:

```
$ ping craggy.hillside.co.uk
craggy.hillside.co.uk is alive
```

You can give an IP address as an argument should you wish. Actually, the command is again hidden away; you'll find it in `/usr/sbin` on Solaris and `/usr/etc` on SunOS. The command needs superuser privilege and is `setuid` to root on all the systems I've looked at.

The `ping` command started life on BSD systems and its default action has been changed somewhat by Sun Microsystems. The original command sent a stream of packets (one per second) until terminated by the user with Control-C. You'll find this original behavior on Linux and BSD systems, and it can be invoked on Solaris or SunOS by supplying the `-s` switch. Listing 2 shows that I'm logged into a machine in California and `ping`ing my system in the United Kingdom.

The message that is sent by `ping` is a standard Internet Control Message Protocol (ICMP) packet. When received by a target machine, it elicits a response that is usually sent from the kernel of the machine's operating system, so you can tell the remote machine is alive. However, the machine may not be supporting users or may be online for other services.

The ICMP message contains a sequence number that is incremented by one for every message that is sent. The returned packet will contain the sequence number, and looking at the increasing sequence numbers can tell you whether or not any packets have been dropped in the round-trip from the sending machine to the remote machine and back. At busy times, you'll find that all networks exercise their right to throw away packets.

Any TCP protocol connection that's carrying important information will cope by retrying to provide a reliable data stream.

The information at the end of each line is the trip time from the sending machine to remote machine and back. It's variable because of network congestion and the other network users. Actually, the average time from California to the United Kingdom two years ago was roughly 250 msec, it's about 180 msec now, reflecting the improved trans-continental links in the United States, and also the faster links used by my ISP across the Atlantic Ocean.

The `ping` command tells you four things. First, it tells you that you can reach the remote machine. Second, that the machine can get packets back to you. Third, that the remote machine is alive, has power and is running some operating system. And finally, it gives you an indication of the speed of the connection between the remote machine and you.

It's a little harder to find out whether or not services are functioning on the remote machine. One approach with some Internet protocols is to use `telnet` to tickle the server. This is viable with FTP, SMTP (mail) and HTTP (Web access). For example,

```
$ telnet machine ftp
```

The `ftp` keyword is looked up in `/etc/services` and translated into a port number. You can always do that yourself and supply a port number as the second argument to the command.

The trick here is to connect and then type `QUIT` to stop the connection. You may also need to know how to crash out of `telnet`: type Control-] to get the `telnet` prompt and then type `close`. If you are bold, before you leave you can also type in some commands from the relevant protocol to make the server do some work. However, do look at the appropriate protocol specification first.

Incidentally, Sun has long provided a command called `mconnect`, which is intended to connect to a mail server and tell you if it's working.

Routes

Sometimes the use of the DNS and `ping` can fool you into thinking that a remote machine is down. I recently had a situation where mail addressed to a customer was hanging around my machine for some time. Simply `ping`ing the machine gave no response, looking up the MX record and the machine name in the DNS showed that the machine didn't exist. Because my

customer had reported poor experiences with his ISP, I suspected that their system had died and the situation was bad. I left it, hoping it would burst into life.

Around 24 hours later, the situation had not altered and further enquiries told me that other people could see the name server and the remote network. When I checked from the United States, I too could see the name server and the network; it looked fine from outside my network, but was invisible from inside.

I used the `traceroute` command to see the routes that packets from my machine took to get to the machine and discovered there was a router problem. My packets were traveling to Germany (rather than staying in the United Kingdom) and were ending up in a black hole. The problem was caused by a routing failure. I contacted my ISP, who sorted things out.

The `traceroute` command was written by Van Jacobsen (of Lawrence Berkeley Laboratories). It's not part of the Solaris release (I haven't managed to look at Solaris 7 yet), but is standard on BSD and Linux. You can find precompiled binaries for Solaris on the Internet. Again, it needs superuser access.

The command makes use of a property of the IP protocol. Each IP packet contains a TTL (time to live) value, which is intended to stop packets traveling endlessly around the world. The TTL is set by the sender and is decremented by every router that the packet passes through. If it reaches one (or zero), then the router is allowed to throw the packet away. However, when doing so, it will send a ICMP message back to the transmitting host saying that it's a time-exceeded packet.

The `traceroute` program uses this behavior to work out the routes packets are taking. First, it sends a packet to the destination host with a TTL of one. The first router will see this, and send back an ICMP error message that can be used to determine its address. Then, `traceroute` sends a message with a TTL of two, getting the next router, and so on until the destination machine is reached. The result is a trace of the route packets have taken to get to the destination and back.

The `traceroute` command allows you to detect loops and black holes, and can sometimes be used to determine where the packets of an unknown IP address emanated from. Beware that some routers block these messages so that you cannot peer into their network.

Further Reading

Much of the basic Internet stuff in this article comes from the definitive, *TCP/IP Illustrated, Volume 1–The Protocols*, by W. Richard Stevens (published by Addison-Wesley Publishing Co., 1994, ISBN 0-201-63346-9). You can get binaries of `dig` and `traceroute` for Solaris from the Solaris Freeware Project at `http://sunfreeware.com`. ❧

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@cpg.com.
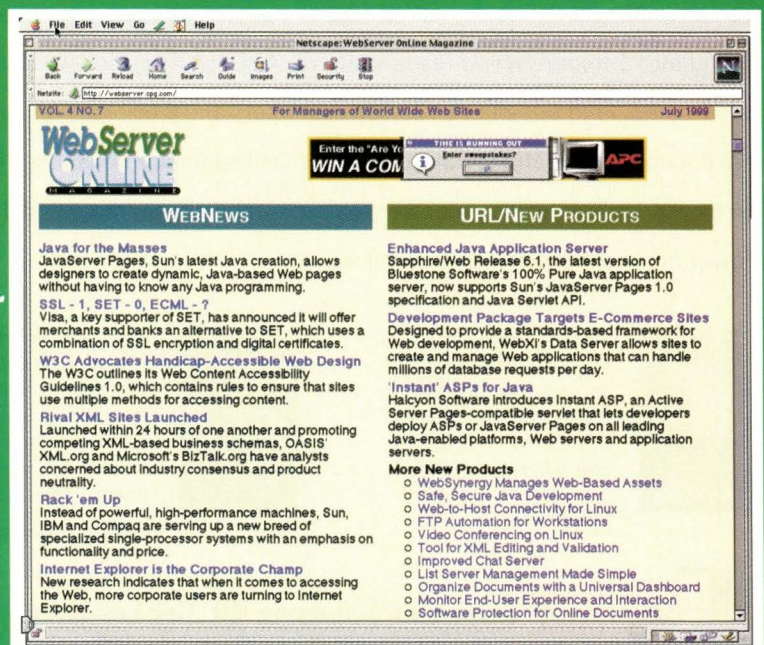
does your **storage** system have enough **FIBRE?**

FIBRE CHANNEL DELIVERS

SPEED · DISTANCE · CAPACITY · FLEXIBILITY

CONSAN IS THE TECHNICAL DISTRIBUTOR OF CHOICE FOR FIBRE SAN SOLUTIONS.

Consan sales reps and technicians have the knowledge and experience to recommend proven Fibre Channel storage configurations for your specific needs.

We carry the top names in Fibre Channel SAN hardware and software such as IBM's family of Ultrastar™ fibre channel hard drives—available in 9, 18 and 36GB capacities. Consan ensures every solution is engineered to meet or exceed your expectations—from simple plug-and-play storage subsystems to custom configured RAIDs and jukeboxes.

To leverage the most of your existing and new storage technologies while increasing speed and data integrity—call Consan—the authority on Fibre SAN Solutions.

1.800.221.6740

IBM®   **CONSAN** A *GATES ARROW*™ Company

18750 Lake Drive East
Chanhassen, MN 55317
TEL: 800-229-3475
FAX: 612-949-0453   www.consan.com

Circle No. 3

NTegration

by Æleen Frisch

LYNNE CANNOY

Upgrade Like it's 1999

By the time this column appears in print, we will have nearly reached the final quarter of 1999. Accordingly, this is the time to face all of those unpleasant Y2K issues that you have been putting off. And if you have not upgraded them recently, your Windows NT 4.0 systems will undoubtedly be among them.

Windows NT Service Pack 5 (SP5) brings Windows NT 4.0 into Y2K compliance. Note that no previous service pack results in complete compliance (not even SP4 followed by the `y2kupd` utility, as Microsoft Corp. previously stated). Detailed information about preparing a Windows NT system for post-2000 dates can be found at `http://www.microsoft.com/technet/year2k/product/user_view69904EN.htm`. This page also lists the compliance status of all available optional add-ons to Windows NT 4.0 (for example, Option Pack components, Distributed File System and so on).

Microsoft's recommended procedure

for upgrading and verifying a Windows NT system comprises two major parts: updating certain facilities not strictly part of Windows NT to their most recent versions–Internet Explorer, FrontPage 97 and the Active Directory Services Interface among them–followed by the SP5 installation. Because virtually all Windows NT systems include Internet Explorer, the procedure for updating it is worth mentioning here.

For Internet Explorer, Y2K compliance comes with Version 4.01 and SP2 (or later). You can determine the version number of a copy of Internet Explorer by selecting the **Help=>About Internet Explorer** menu path and examining the Version and Update Version fields displayed, ensuring the latter includes "SP2" somewhere within the string. You can download Internet Explorer 4 or 5 from `http://www.microsoft.com/windows/ie/download/all.htm?bShowPage`.

As with other Windows NT service packs, SP5 contains all of the operating

system updates and cumulative fixes provided in service packs 1 through 4 (and accordingly, it is not necessary to install any previous service packs before applying this one). In this way, Windows NT service packs represent what other vendors typically refer to as minor operating system revisions. One aspect about SP5 that is unusual is the fact that it introduces no new functionality to the operating system; it merely corrects problems introduced by SP4. As such, SP5 is really a fixed version of SP4.

The normal location for obtaining Windows NT service packs is in the subdirectories of `ftp://ftp.microsoft.com/bussys/winnt/winnt-public/fixes/usa/nt40` (users outside the United States should replace "usa" in the preceding path with the appropriate country code). However, for reasons which have not been specified, SP5 is delivered from an alternate location: `http://www.microsoft.com/ntserver/nts/downloads/recommended/sp5/allsp5.asp`.

From here, you can download the service pack, order a CD version of it, view its `readme` file and perform other similar functions. Checking out the `readme` file before installing the service pack is important in order to be made aware of any hardware incompatibility and/or software support issues (especially for third-party software). Note that service pack contents do not vary according to whether they will be installed on a server or workstation (despite the naming of the current download directory).

When you're ready to download the software, select the version that is appropriate for your hardware (Intel vs. Alpha) and security configuration (40-bit vs. 128-bit). U.S. users can find out if their system is using 128-bit security by determining if a file named `rsaenh.dll` exists on the system disk. Once you have made these selections, the next dialog asks you to specify the appropriate language for the update.

Finally, you are asked to choose between an "Express" download and a "Full" download procedure. Express downloads are designed for applying software to the local system only. In this mode, a setup program is downloaded and then immediately executed. This program queries the local system to determine its hardware and software configurations and then downloads only those files from the service pack that are needed. Once this second download operation is completed, service pack installation will begin automatically.

In contrast, a Full download operation retrieves the entire service pack (packaged as a single executable file), which can then be applied to numerous systems, regardless of their specific configurations (because every possible file is included). This executable may be run by double-clicking on its icon or by entering its name at any command prompt. Service pack installation will then begin as the executable itself is verified and unpacked. Next, you'll be prompted to accept the accompanying license agreement and asked whether you wish to create an uninstall directory, which will enable you to remove the service pack later if necessary. This dialog is illustrated in Figure 1 on Page 38.

Figure 1. Specifying the Creation of an Uninstall Directory

Windows NT Service Pack Setup

Welcome to Windows NT 4.0 Service Pack 5 Setup.

Before installing this Service Pack, we recommend that you close all other applications, backup your system, and update your Emergency Repair diskette. When Service Pack Setup completes, you will need to shutdown and restart Windows NT.

To continue, please read the following license agreement and indicate your acceptance:

SUPPLEMENTAL END USER LICENSE AGREEMENT FOR MICROSOFT SOFTWARE

IMPORTANT: READ CAREFULLY - The Microsoft Corporation ("Microsoft") operating system components which you are about to install, including any "online" or electronic documentation ("OS COMPONENTS"), are subject to the terms and

☑ Accept the License Agreement (must accept before installing the Service Pack)

☑ Backup files necessary to uninstall this Service Pack at a later time (requires

Figure 2. Handling Vendor-Provided Files

Windows NT Service Pack Setup

C:\WINNT\System32\DRIVERS\E100B.SYS

Your computer vendor installed this file on your computer. Do you want Service Pack Setup to replace this file?

[Yes] [No]

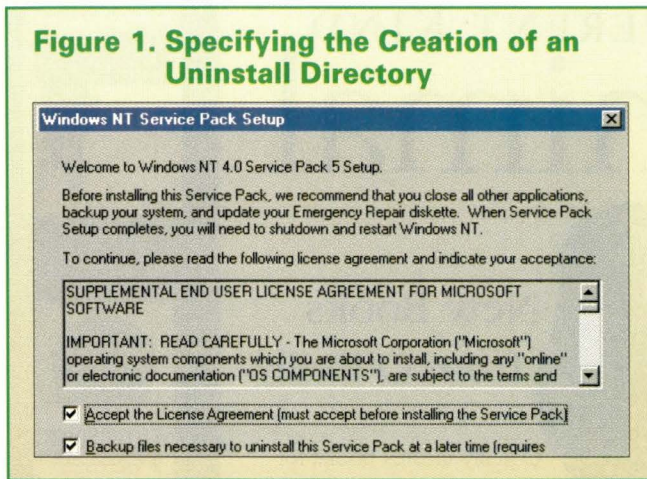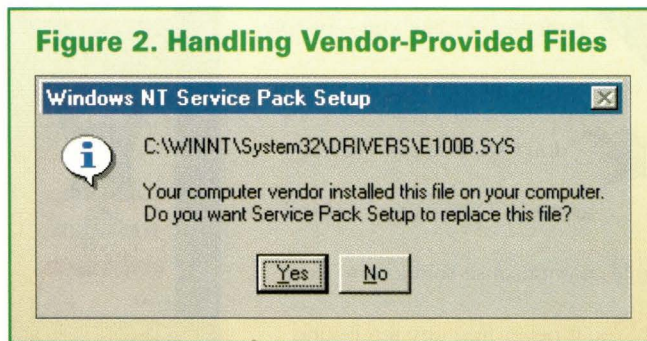If you have sufficient disk space, creating an uninstall directory is a prudent idea. Once you have made these choices, installation of the actual service pack files then proceeds. At some point in the process, you may see a dialog like the one shown in Figure 2. This message indicates that the version of the specified file currently on the system was provided by a non-Microsoft vendor (usually a hardware manufacturer). The installation program informs you of this and allows you to choose which version of the file to use (rather than automatically replacing it with the version included in the service pack). In my experience, I have found it is better to retain the vendor's version rather than accepting the one from the service pack. Once all of the service pack files have been installed, the system will usually reboot.

One final step remains when the system returns. In order to use an emergency repair disk in the future to repair a damaged system, one of the Windows NT installation disks must be modified slightly (as noted in the service pack's release notes). Specifically, you will need to copy the version of the setupdd.sys file included with the service pack to Install Disk 2, replacing the version found on the original diskette (this step has been necessary since SP2). Unfortunately, this file is not included in all of the downloadable versions of SP4 and SP5. However, you may obtain it by downloading the architecture-specific executable file located at ftp://ftp.microsoft.com/bussys/winnt/winnt-public/fixes/usa/nt40/ussp4/Additional, and then running it from a command line with the /X option. For example, sp4exti.exe /x for the Intel version.

Service packs can be reinstalled on top of themselves with-out causing problems, and must be reinstalled whenever new hardware, system services or facilities are added to the system. If you decide that you want to remove a service pack, you can do so using the files stored in the %SystemRoot%\WinNT\$NTServicePackUninstall$ directory. Specifically, run the spuninst command found in the spuninst subdirectory at this location. Once you have decided to keep the service pack, you may delete this entire subtree to save disk space.

Are We There Yet?

If you completed the procedures described in the preceding section (as well as any other applicable steps documented on the Y2K information Web page), your system should be ready for the new millennium. However, no service pack installation should be regarded as the last word in operating system updates. Between service pack releases, Microsoft makes available what it calls "hot fixes," which are patches to a specific part of a given Windows NT release. As such, they are generally designed to fix bugs discovered since the release of the most recent service pack (and sometimes even introduced by it).

Since SP5, several new Windows NT security problems have been found. For example, buffer overflow bugs in both the Remote Acess Service (RAS) and Help File facilities can allow an unscrupulous user to execute "arbitrary code," allowing him to gain administrative access to the system. Fixes for these problems are available from ftp://ftp.microsoft.com/bussys/winnt/winnt-public/fixes/usa/nt40/Hotfixes-PostSP5. The fix for the Help File problem, for example, is located in the winhlp32_fix subdirectory. To install the fix, download and run the appropriate architecture-specific executable file from this location.

An excellent resource for information about these and other security-related topics is www.ntsecurity.net (please note that www.ntsecurity.com is a different site). As of this writing, the site provides information on all post-SP5 security issues, links to the corresponding Microsoft-provided documentation and hot fixes, as well additional information about a serious security problem found in Internet Explorer (see http://www.microsoft.com/windows/IE/security/favorites.asp for fixes) and one as-yet unresolved Windows NT security vulnerability (see "NT Case Sensitivity Vulnerability" at www.ntsecurity.net for details).

As noted earlier, in addition to fixing problems, service packs usually provide expanded functionality and additional utilities to the basic Windows NT operating system. SP4 did so, and these new features also carry over into SP5. Next time, we will look at some of the most interesting and important of these new features. ➟

Æleen 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.

Work

by Jeffreys Copeland and Haemer

"Those who cannot remember the past are condemned to repeat it."
– George Santayana

"Those who do not understand UNIX are condemned to reinvent it, poorly."
– Henry Spencer

Software Ptools

This month, we take you to Hershey Heaven. In the flurry of irrelevancy with which the press inundates Nobel Prize winners, Alfred D. Hershey, the 1969 Nobel Laureate in Physiology or Medicine, was asked what he thought heaven would be like. He answered that in heaven, he thought he'd finally get an experiment that worked–and be able to do it over and over again.

In 1976, Brian W. Kernighan and P.J. Plauger wrote *Software Tools* (published by Addison-Wesley Publishing Co., ISBN 0-201-03669-X), a book we recommend to everyone. Here's a brief synopsis supplied by coauthor Peter Plauger:

Describes a number of small programs made popular by the UNIX operating system. Contains complete source code of all the programs in Ratfor, a structured dialect of FORTRAN that strongly resembles C. This classic pioneered the term 'software tools.'

This book, probably the clearest exposition of the UNIX tools philosophy, provides the source code for complete implementations of several UNIX tools, together with running commentary on every point of their design and implementation.

But why in Ratfor–something that "strongly resembles C"–and not C itself? The first edition of Kernighan and Ritchie's *The C Programming Language*, which introduced the world to C, was published in 1978. But in 1976, the two universally available programming languages were COBOL and FORTRAN-66 (often as the implementation FORTRAN IV). Kernighan and Plauger had a message to get out about how to program, to an audience who had never heard of UNIX or C, and who–at least the way it looked then–never would.

As a transition vehicle, they created a language that looked like C, but could be preprocessed into FORTRAN-66. (Most of you have probably never seen FORTRAN-66, but it lacks nearly everything you take for granted in programming languages: if-then-else, data structures, while loops, character I/O and even strings. The primary control-flow structures are the "logical if" and the "goto.")

Here's an example of relatively easy-to-read FORTRAN-66, taken from the original Bell Labs Ratfor documentation:

```
IF (X.LE.100) GOTO 10
   CALL ERROR(5HX>100)
   ERR = 1
   RETURN
10 ...
```

The equivalent in Ratfor?

```
if (X>100) {
   call error("x>100"); err = 1;
   return }
```

The idea was to use a language that would permit easy-to-read examples in a book that showed folks how to write code that improved their programming

environments. To close the loop, the final chapter of *Software Tools* designs and implements an entire Ratfor-to-FORTRAN preprocessor.

The book quickly spurred the formation of the Software Tools User Group (STUG). Formed at Lawrence Berkeley Labs, this group began distributing tapes that contained all of the code from the book, together with an ever-growing body of contributed tools, all of which could be installed on any computer with a FORTRAN compiler (which meant, at the time, pretty much any computer)–a UNIX-like environment, not just for non-UNIX systems, but for a world that had never heard of UNIX or C.

If you'd like to see what sorts of things were done, you can still find references to *Software Tools* on the Web. One such site is `http://www.geocities.com/SiliconValley/Lab/9247/#compilers`. An advantage of the UNIX "one tool, one job" philosophy is that you can attack each command separately, one at a time. Most are bite-size. Individual contributors can write a useful tool in a few days (or less) and make a real contribution to the larger whole. Yoked to this is the idea that when you write a little program, instead of a giant, monolithic one, you can really make it yours; you can get your arms around it and put in the work you need to get it just right.

The reason we can recommend a 25-year-old book, full of code for programs you'll never need to write, in a language you'll never use, is that it remains the clearest, best written, most entertaining and practical treatise we know on how to get programs just right. (Do not, by the way, be fooled into buying *Software Tools in Pascal*, written by the same authors and published in 1981. Kernighan's 1981 technical report, entitled, "Why Pascal Is Not My Favorite Programming Language," `http://cm.bell-labs.com/cm/cs/cstr/100.ps.gz`, gives great insights into both why this book didn't turn out the way it could have and why Pascal, once a hot contender for the programming language of choice, eventually lost out to C.)

It's hard to imagine, nowadays, just how revolutionary this book's approach was. It changed lives. Most folks reading this column have probably never even seen a punched card or written a FORTRAN program. To help put things in context, imagine working as a programmer in a world in which neither you nor anyone you know has ever heard of a filter or a "software tool," and the only tools available to you as a programmer are a compiler, an assembler and a linker. (And a world where nothing is off-the-shelf. We know someone who began his career writing a payroll system in FORTRAN for a movie studio.)

One chapter of *Software Tools* contains the complete design and implementation of an editor. Not a screen editor, mind you. After all, no one had cursor-addressable terminals back then.

What ever happened to STUG? One finds occasional fossils of STUG, such as Usenix's Software Tools User Group Award, `http://www.usenix.org/directory/stug.html`, but the group died of its own success. People who joined STUG learned about UNIX, helped popularize UNIX and the UNIX philosophy, eventually demanded UNIX and switched to UNIX when it became available.

In the mid-1980s, the UNIX tool set story was replayed more than once, to the great advantage of a new generation of computer users. Mortice Kern Systems (MKS) Inc., a Canadian company, rewrote the entire basic UNIX command set from scratch for MS-DOS, and later ported the same suite to various legacy systems.

In the same time frame, the Free Software Foundation (FSF) coordinated the contribution of an army of volunteers, who created freely redistributable versions of nearly all common UNIX tools, which eventually made up the bulk of the command-line utilities for Linux.

A third great source of rewritten UNIX tools are the BSD releases, coordinated by the University of California at Berkeley's Computer Science Research Group (CSRG), and found in a wide variety of freely available BSD-based UNIXes. As with STUG and the FSF, CSRG's work was the coordinated effort of an unruly army of individual volunteers.

Everyone eventually caught on to the software toolbox approach.

Tom Christiansen Becomes Irked

Well, everyone except the Windows world.

Those of you who read the `comp.lang.perl.misc` newsgroup know that Tom Christiansen, coauthor of many O'Reilly & Associates Inc. Perl books, is a frequent contributor. When a question irks Tom, he speaks right up, often chastising the questioner. Some people don't like this, so Tom is occasionally a source of discussion on the newsgroup in his own right.

Like grains of sand in an oyster, though, these irritants sometimes spur Tom to create something beautiful. (In the distance, we hear groaning; to quote Jo Haemer, "A cheap shot is a terrible thing to waste.")

Spurred on by irksome questions, Tom has written Perl man pages, FAQs, tools and even a series of Perl FMTYEWTK (Far More Than You Ever Wanted To Know) essays.

Several months ago, Tom was going through a stretch of irritation at people asking for complete Perl solutions to problems that could be solved with simple calls to basic UNIX utilities. Too often, his pointing this out didn't help the requester, because they were running some Microsoft Corp. platform that didn't have the basic utilities to solve the problem.

For a while, Tom's reaction was to declare that such lacunae were God's wrath visited upon anyone sinful enough to run something other than UNIX, and that we in the UNIX community had no obligation to help.

Those of you who've been reading the newsgroup, or this column, for some time will even remember a parody posting, by Nat Torkington–Tom's coauthor for *The Perl Cookbook* (published by O'Reilly, 1998, ISBN 1-56592-243-3).

Tim.Bunce@ig.co.uk (Tim Bunce) writes:
> The problem is to find the full list of
> names and the original order.

You INSTALL a FULL SET OF TOOLS, like THE LORD GOD
ALMIGHTY intended. REPENT, ye PRISONER of BILL!
The DAY of JUDGEMENT is AT PERL! Your MESSENGERS
are obviously just POOR substitutes for RELIABLE
PIPE COMMUNICATION which you'd have if you had a
REAL OPERATING SYSTEM and not a SCURRILOUS PIECE
OF TOOL-CHALLENGED COPROPHILIA!

Tom^WNat
:-)

This ultimately led us to write a Perl version of tsort(1) (see http://sw.expert.com/C9/SE.C9.SEP.98.pdf and http://sw.expert.com/C9/SE.C9.OCT.98.pdf).

More recently, though, Tom seems to have decided that the problem isn't going to go away, and has organized a project to rewrite all the basic UNIX utilities in Perl, so that any system with Perl can have the full, basic UNIX command set for free.

Note the word "organized." Tom has written some utilities himself but what he's really doing is coordinating contributions from all over the Perl world. Tom calls it the "Perl Power Tools" project. We prefer "Software Ptools": the "P" is psilent.

An Example: asa(1)

This looked like enormous fun, and we jumped in with both feet. We have both worked in the printer industry, so we decided to chip in by contributing a traditional UNIX utility that no one else would be silly enough to write: asa(1), a program that interprets traditional FORTRAN carriage-control commands.

Here, a little history will help. A couple of decades ago, printers were all impact line printers that produced great stacks of accordion-folded, green-and-white-lined,14-inch-wide paper, which was actually 15-inches wide if you counted the tear-off strips on the sides that were perforated so teeth on the printer carriage could advance the paper.

These printers were simple; most couldn't even do graphics (graphics output was provided by another kind of printing device, called a plotter). Indeed, besides printing alphabetic characters, almost all they could do was to move to the top of a new page, backspace (in order to underline) and overprint lines (to produce bold characters).

But at the time, not even character sets were portable (non-IBM Corp. machines often used ASCII, but IBM machines, which were in the majority, used EBCDIC), which meant your programs couldn't assume that backspace was a ^H, form-feed was a ^F or carriage return was a ^M.

A convention was born: all printers agreed to look at the first character of each output line and interpret a small number of special characters as special carriage-control commands. For example, a "1" in the first column of an output line told the printer to eject the current page and move to the top of a new page. These conventions were made a part of the American Standards Association (ASA) FORTRAN standard. (You read that correctly: printer carriage controls were part of a programming language standard.) ASA was later renamed ANSI.

In the C/UNIX world, there are no such conventions. Moreover, neither contemporary terminals nor newer printers, such as laser printers, interpret output in this way. Old FORTRAN programs, ported from other operating systems, began finding themselves assuming these conventions on systems that didn't recognize them.

To handle this, early UNIX systems included a program called asa(1), which translated FORTRAN carriage controls. Listing 1 shows our implementation of asa(1).

And now for our dramatic reading: The meat of the program lies in 10 lines of code, lines 10 through 20. Everything else is professionalism.

Listing 1. asa(1)

```
1   #!/usr/local/bin/perl -w
2   # $ID: asa,v 1.1 1999/05/31 22:03:15 jsh Exp jsh $

3   use strict;

4   exit 1 if grep {!-r} @ARGV;    # traditional

5   if (grep /-/, @ARGV) {
6     $0 =~ s(.*/)();
7     warn "usage: $0 [filename ...]\n";
8     exit 2;   # traditional
9   }

10  while (<>) {
11    chomp;
12    s/^$/ /;
13    s/^[^10+-]/\n/;
14    s/^1/\f/;
15    s/^\+/\r/;
16    s/^0/\n\n/;
17    s/^-/\n\n\n/;
18    print
```

```
19      or exit 1;    # traditional
20   }

21   =head1 NAME

22   asa - interpret ASA/FORTRAN carriage-controls

23   =head1 SYNOPSIS

24   asa [I<filename> ...]

25   =head1 DESCRIPTION

26   =over 2

27   Traditional FORTRAN programs put carriage-control characters
28   in the first columns of their output,
29   which were interpreted by older line printers
30   according to the ASA vertical format control standard.
31   (ASA was the American Standards Association -- now ANSI.)

32   Under this standard, the first character of each printable record (line)
33   determines vertical spacing, as follows:

34   =over 2
35   I<blank>     carriage return
36   0            two carriage returns
37   1            formfeed
38   +            overprint
39   -            three carriage returns (IBM extension)

40   =back

41   All other characters are discarded, and empty lines behave as though
42   they have a leading blank.

43   B<asa> interprets these characters.

44   =back

45   =head1 EXIT VALUES

46    =over 2

47   0 normal exit

48   1 inability to write on stdout or to read an input file

49   2 bad argument

50   Exit status values chosen from MKS toolkit.

51   =back

52   =head1 AUTHOR

53   Jeffrey S. Haemer

54   =head1 BUGS

55   Currently, B<asa> just looks at the readability of its input files
56   at startup time. It should really do it a file at a time,
57   but that makes the code look gross.

58   The carriage-control '-' is an IBM extension.
59   Perhaps the default should ignore it
60   and there should be a '-i' option to interpret it.

61   =head1 SEE ALSO
62   I<Communications of the ACM>, Vol 7, No. 10,
63   p. 606, October 1964.

64   NWG/RFC 189, Appendix C

65   =cut
```

Lines 1 through 3 are our usual boilerplate. The shebang line (line 1) invokes the Perl interpreter and gives it the -w flag, which queries various questionable usages. The third line requires the still more picky strict pragma. As long as we're going to write a utility, we might as well catch as many silly errors as we can. The second line says we're keeping our code under revision control.

Lines 4 through 9 do argument parsing. The comment "traditional" means that it's traditional for this command to exit with an exit status of "2" if the arguments are misspecified.

Lines 21 through 65 are documentation. Perl lets you keep your documentation in the same file as your code, so they don't get out of sync.

The meat of the program is the loop begun on line 10 and finished on line 20, which reads and prints the file one line at a time. Carriage control is specified entirely by the first character in the line, so line 11 begins by removing any special ASCII carriage controls at the ends of lines. The printer will never see them. The standard says that a line beginning with anything except one of the special ASA carriage-control characters should trigger a new line and a carriage return, and lines 12 and 13 give us that. Line 12 prints blank lines as blank lines. Line 13 consumes any other character that begins a line and performs the default action: terminating the preceding line. (Yes, that's right. If you want a character at the beginning of a line, you have to precede it with something else. The first character is always interpreted as carriage control.)

Lines 14 through 17 interpret the ASA, beginning-of-line carriage-control codes:

1	Form-feed
+	One carriage return (for overprinting)
0	Two carriage returns/line-feeds (for double-spaced lines)
−	Three carriage returns/line-feeds (for triple-spaced lines)

There it is. A 10-line program. Just an oddly shaped brick in the ziggurat of free UNIX tools for the non-UNIX world.

Want to chip in? It's fun. Go to http://language.perl.com/ppt and take a look at what's done and what's not. Make the world a better place by spending a few hours in Hershey Heaven.

Until next time, 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 *or alternately at* ftp://ftp.expert.com/pub/Work.

MILAN KECMAN

Across the Great Divide

Originally, computer networks were intended to allow people spread out over a wide geographical area to access large computers stored at central locations. At the time, computers were extremely expensive and were only cost-effective if many people could use them at once. Over time, computers have shrunk to the point where it has become cost-effective for everyone to have their own computer, one that is dedicated to doing the particular tasks the individual requires. Unfortunately, this progression led to the isolation of information (information needed to be copied to individual computers to be used by more than one person), as well as the requirement that each computer have a complete set of peripherals even though they may be rarely used.

As with the large centralized servers, it often makes sense to share computing resources among multiple users to avoid unnecessary duplication or to achieve a better economy of scale. Computer networks have now evolved to the point

where they are used to share resources such as disk drives and printers among several computers, thus providing an augmented computing environment more powerful than could be afforded by most individuals and eliminating the unnecessary duplication of information.

As networks have evolved, so too has the software that is used to communicate between them. Originally, software that talked across networks was very simple; usually providing either a terminal interface or some simple data transfer operation (email, file transfer and so on). As local computational power has become less expensive, however, it has become possible for more and more tasks to take place on the local machine rather than on the central machine. This has led to an explosion in task-specific network protocols, and many of them have very fine data transfer granularities.

One problem with network software is that getting applications to talk coherently is painstakingly difficult. Software wishing to use services running on

another computer must get the other computer's attention, ask it questions and interpret the responses. Similarly, software wishing to provide services to other computers must make its presence known to those computers, watch for incoming questions, interpret them, perform a task, formulate a response and send it back. While not conceptually difficult to manage, there is so much housekeeping involved in this process that writing error-free network-aware applications has become extremely difficult and expensive.

Because many of the housekeeping procedures are common among network-aware software, the obvious thing to do is build a package able to describe the data you want to send and the data you expect to receive in a convenient, readable format and have a software package build you the code that handles all the details.

This same problem was encountered years ago as computer programs became more complex. Programmers spent so much time doing the same mundane

Listing 1. RMI Interfaces

```
/** An RMI interface that provides file finder services. */
public interface RemoteFileFinder extends Remote
{

    /** The name of the RMI service. */
    final static String SERVICE_NAME = "FileFinder";

    /** Opens up a file on the remote file server for read-access */
    RemoteFileReader openForRead(String fileName)
        throws IOException, RemoteException;
}

/** An interface that allows a remote file to be read. */
public interface RemoteFileReader extends Remote
{

    /** Reads up to the specified number of bytes from the remote file.
     * Returns null if no more data. */
    byte[] read(int size) throws IOException, RemoteException;

    /** Closes the remote file. */
    void close() throws RemoteException;
}
```

Listing 2. The RemoteFileFinder Implementation

```
/** A network file server based on RMI. */
public class RemoteFileFinderImpl extends UnicastRemoteObject implements RemoteFileFinder
{

    public RemoteFileFinderImpl() throws RemoteException {}

    /** Opens up a file on the remote file server for read-access */
    public RemoteFileReader openForRead(String fileName)
        throws IOException, RemoteException
    { return new RemoteFileReaderImpl(fileName); }
}
```

tasks (and making the same inevitable mistakes while doing it) that the obvious solution was to create a new language in which you could describe what you wanted to do, in convenient little pieces, and have a computer program translate that description into native computer code. Those programs are today's compilers and the pieces are methods. Methods talk among themselves via method calls, which are standardized techniques for passing information back and forth between code segments.

If this technique could be applied to a single application, why couldn't it be used to produce programs that operate across a network? And why couldn't it be done in such a way that the application thinks that it is just making a normal method call?

Java provides this capability via the *remote method invocation* (RMI) system.

The Parts and the Whole

Every RMI-based application is composed of three parts: the *server*, which exports objects to one or more *clients*, and a *registry*, which allows the clients to obtain references to these objects by looking them up by name.

To create an RMI server, the first thing you need to do is create an interface describing the services that you would like to provide to clients. RMI object interfaces have two special requirements. First, the interface must extend `java.rmi.Remote`, which marks any class implementing the interface as RMI-enabled. Second, all methods provided by the interface must throw `java.rmi.RemoteException`, in addition to any exceptions they might otherwise throw. Our example application, which may be downloaded in its entirety from `ftp://ftp.expert.com/pub/JavaClass/08.1999/rmiserver.tar`, provides two RMI interfaces: RemoteFileFinder and RemoteFileReader (see Listing 1).

Once you have created the RMI interfaces, you must create a class that implements them. This class will provide the server-side implementation of the object. As is the case with the RMI interfaces, these classes have two special requirements: they must extend one of the subclasses of `java.rmi.server.RemoteServer` and must throw `java.rmi.RemoteException`, in addition to any exceptions they might otherwise throw. Normally, RMI classes extend `java.rmi.server.UnicastRemoteObject`, the only such subclass provided by Java Development Kit (JDK) 1.1. Our example application's RMI interfaces are implemented by Remote

Listing 3. The RemoteFileReader Implementation

```java
/** Implementation of the RMI remote file reader interface. */
public class RemoteFileReaderImpl extends UnicastRemoteObject implements RemoteFileReader

{
   /** File input stream used to read file data. */
   private FileInputStream inStream;

   /** True if the file data has been completely read. */
   private boolean atEof;

   /** Output stream used to accumulate file data */
   private ByteArrayOutputStream outStream = new ByteArrayOutputStream();

   /** Creates a remote-able file reader object. */
   public RemoteFileReaderImpl(String fileName) throws IOException, RemoteException
   { inStream = new FileInputStream(fileName); }

   ///////////////////////////////////////////////
   // RemoteFileReader interface implementation //
   ///////////////////////////////////////////////

   /** Reads an array of bytes from the remote file.  Returns null if there is no more data. */
   public byte[] read(int size) throws IOException, RemoteException
   {
      if (atEof)
         return null;
      outStream.reset(); /* clear output stream */
      while (size >= 0) {
         int value = inStream.read();
         if (value < 0) {
            atEof = true;
            break;
         }
         outStream.write(value);
      }
      return outStream.toByteArray();
   }

   /** Closes the file. */
   public void close() throws RemoteException
   {
      try {
         inStream.close();
      }
      catch (IOException e) {}
      atEof = true;
      outStream = null;
      inStream = null;
   }
}
```

FileFinderImpl and RemoteFileReaderImpl (see Listing 2 and Listing 3, respectively).

In an RMI application, the client manipulates server objects via a proxy object, called a *stub*. You do not have to write these proxy objects; the rmic compiler provided with JDK 1.1 does it for you. Supply rmic with the names of the server-side class(es), as follows:

```
% rmic fileserver.RemoteFileFinderImpl
```

This creates two new classes: fileserver.RemoteFile FinderImpl_Skel and fileserver.RemoteFile FinderImpl_Stub. The former performs server-side manipulations, while the latter is the proxy object used by the client. Normally, you need not concern yourself with these classes; they are used transparently by the RMI framework.

Hear Me Roar

Once a server has created an instance of an RMI-enabled interface, it needs to tell the registry about it so that a client can find it. The JDK provides a registry service, *rmiregistry*, which may be used to provide a common registry for RMI objects across many Java Virtual Machines (JVMs). In practice, however, it is often confusing to the administrator to have to run both the registry service and the RMI application, particularly because the registry service must be started before the RMI application. For this reason, it simplifies matters to have each RMI application create its own registry with the java.rmi. registry.LocateRegistry.createRegistry(int) method, which creates an RMI registry on the indicated port. This is the first task of the Server class in our sample application (see Listing 4). The application can then bind an RMI name to an RMI object with the bind() or rebind() methods, after

Listing 4. The Server Application

```java
/** Class which creates our file server and registers it with an RMI registry. */
public class Server
{
    /** Port number for RMI registry */
    public static final int SERVER_PORT = 10001;

    /** Starts up the RemoteFileFinder RMI service. */
    public static void main(String[] args)
    {
        // start an RMI registry for this server
        Registry registry;
        try {
            registry = LocateRegistry.createRegistry(SERVER_PORT);
        }
        catch (RemoteException e) {
            System.err.println("Error while creating RMI registry:");
            e.printStackTrace();
            return;
        }

        RemoteFileFinder finder;
        try {
            finder = new RemoteFileFinderImpl();
        }
        catch (RemoteException e) {
            System.err.println("Error while creating remote file finder:");
            e.printStackTrace();
            return;
        }
        // register the file finder object with the RMI registry
        try {
            registry.rebind(RemoteFileFinder.SERVICE_NAME, finder);
        }
        catch (Exception e) {
            System.err.println("Error while binding RMI service:");
            e.printStackTrace();
        }
        System.out.println("File server started up at rmi://localhost:" + SERVER_PORT + "/" +
                        RemoteFileFinder.SERVICE_NAME);
    }
}
```

which they may be found and used by RMI clients.

The `java.rmi.Naming.lookup(String)` method provides the ability to resolve a URL for an RMI object, for example, `rmi://host_name:registry_port/service_name`, into an RMI object stub. This is the first task performed by the Client class in our example application (see Listing 5, Page 49). Once you have the object stub you may use it as if it were a normal Java object. To run the example server, type (command is wrapped for printing)

```
% java fileserver.Server
File server started up at
   rmi://localhost:10001/FileFinder
```

To connect to the server and display a file's contents, type

```
% java fileserver.Client
   rmi://localhost:10001/FileFinder /etc/motd
Welcome to Snowplow
```

In just a few lines of code we've created an RMI service that performs much of the task of that of a file transfer program.

The Man Behind the Curtain

So how does a remote method invocation actually work? When the client calls a method on an RMI object, reference control is passed to the stub class that was generated by `rmic`. This class takes the arguments to the method and uses the Java serialization system (see `http://www.java.sun.com/products/jdk/1.1/docs/guide/serialization/index.html`) to pass them off to the server via an IPC mechanism (currently an Internet socket). The Skeleton class, also generated by `rmic`, reconstitutes the objects and calls the appropriate method on the server object. If the method returns normally, its return value is serialized and passed back to the client stub, which reconstitutes it and returns the value to the application. If it throws an exception, the exception is serialized (with some loss of information), passed back and thrown on the client side.

Because of the network layer between the client and server, there is an additional level of unreliability that is not seen in a normal application. This is why the RMI interfaces must throw `java.rmi.RemoteException`; if an error (such as a network failure) occurs while RMI is processing the request, this exception will be thrown and the client can attempt to recover.

There is one important case in which an argument or return

Java Class

Listing 5. The Client Application

```
/** Test client for the RMI file server. */
public class Client
{

    public static void main(String[] args)
    {
        // Find the file finder using the server's RMI registry
        RemoteFileFinder finder;
        try {
            finder = (RemoteFileFinder)Naming.lookup(args[0]);
        }
        catch (Exception e) {
            System.err.println("Exception while trying to connect:");
            e.printStackTrace();
            return;
        }

        // try to open the file
        RemoteFileReader reader;
        try {
            reader = finder.openForRead(args[1]);
        }
        catch (FileNotFoundException e) {
            System.err.println(args[1] + ": File not found");
            return;
        }
        catch (Exception e) {
            System.err.println("Exception while trying to open file:");
            e.printStackTrace();
            return;
        }

        // read the file contents
        for (;;) {
            try {
                byte[] buffer = reader.read(1024);
                if (buffer == null)
                    break;
                System.out.print(new String(buffer));
            }
            catch (RemoteException e) {
                System.err.println("Exception while trying to read file data:");
                e.printStackTrace();
                break;
            }
            catch (IOException e) {
                System.err.println("I/O error while reading file data:");
                e.printStackTrace();
            }
        }

        // clean up
        try {
            reader.close();
        }
        catch (RemoteException e) {
            System.err.println("Exception while cleaning up:");
            e.printStackTrace();
        }
    }
}
```

value is not serialized when passed between the client and the server—that is, when the object implements the `java.rmi.Remote` interface. In this case, an object reference is passed in place of the object value, telling the RMI system that this object should be manipulated via a proxy stub rather than reconstituted locally. This feature is very useful; it means servers can be written using many fine-grained interfaces and without a lot of programmer effort, allowing servers to be much more extensible. This feature is used by the `RemoteFileFinder.open ForRead()` method in our example application.

Next Month

So far in this column we've looked at many of Java's built-in features, which are very powerful. In some cases, however, Java does not provide the capabilities required by an application. For such cases, Java provides the Java Native Interface, or JNI, which will be the topic of next month's column. ➙

Jim Frost is a software engineer specializing in Java technologies and strong opinions. He may be reached by email at jimf@frostbytes.com.

Software Development in an Object World

Visual this, visual that, ASP, EJB, UML, outsourcing—sorting out development options has become a real puzzle.

With Comdex weighing in with more than 125,000 attendees and vendor-specific annual conferences from the likes of Computer Associates International Inc., Hewlett-Packard Co., Microsoft Corp. and Oracle Corp. routinely attracting more than 10,000 developers, Software Development '99 West (*http://www.sdexpo.com*) certainly isn't the largest conference on Earth. However, it is arguably the best general-purpose software development conference out there and has always attracted a large number of vendors and corporate attendees. Its sessions reliably mirror the software development community's concerns and trends.

This year's conference was spread out over five days in May and included a management track, in addition to nine content-rich developer tracks (C++ programming, database and middleware development, D/COM+ programming, Java programming, Internet development, methods and modeling,

by Karen Watterson

object and component development, user interface design and Windows 2000 platform development). According to the show's sponsor, the most popular sessions were Scott Meyers' "Real World C++" and "New Cool Things in C++," Martin Fowler's session on "Analysis Patterns" and Terry Quatrani's "Introduction to Visual Modeling and UML."

That this show would feature tracks on C++, Java, the Internet and Windows programming shouldn't come as a surprise; but what can the popularity of sessions on UML and analysis patterns tell us?

UML and Patterns

UML, which stands for Unified Modeling Language and was pioneered by Rational Software Corp., Cupertino, CA, is a standard for specifying, visualizing, constructing and documenting the artifacts of software systems. UML was officially adopted by the Object Management Group (OMG) and is slowly emerging as the software modeling equivalent to SQL, the interoperable Structured Query Language standard used by database vendors. Most of today's computer-aided software engineering (CASE) tools support UML, and there are scores of books on the subject. Patterns are also a popular topic today and reflect the interest of developers in the reuse of components that are much larger-scale than, say, ActiveX or JavaBeans components.

"A pattern," says Fowler in his book *Analysis Patterns* (published by Addison-Wesley Publishing Co., 1996, ISBN 0-201-89542-0), "is an idea that has been useful in one practical context and will probably be useful in others." The notion of patterns was first popularized by the so-called "gang of four" (Erich Gamma, Richard Helm, Ralph Johnson and John Vlissides) in the classic *Design Patterns* (also published by Addison-Wesley, 1995, ISBN 0-201-63361-2). Patterns in general have become an important part of object-oriented development.

Analysis patterns basically represent reusable object models and encapsulate domain knowledge in areas as diverse as accounting, customer service and health care. Savvy developers realize that by understanding design and analysis patterns, they can acquire some of the business savvy that most IT shops expect today. It's no longer enough to merely be a good C++ or Visual Basic programmer—you're expected to understand something about business problems as well.

Speaking of CASE, remember "upper" and "lower" CASE? Upper CASE tools referred to those high-end soup-to-nuts packages that were often part of a methodology. (Methodologies were essentially detailed recipes for building complex systems.) Lower CASE tools were designed to solve a particular modeling problem, notably, database modeling. Have CASE tools and the methodologies of the late '80s and early '90s evolved to support today's object-oriented development environment?

Interestingly, there wasn't a single session on CASE (although there was one on the object-oriented technique of "use cases") and only one on methodologies. That's not to say the basic software development issues CASE and methodologies tackled have disappeared. You might say UML is today's answer to CASE, and that patterns, knowledge management (think of the vast databases of "best practices" maintained by the large consulting firms) and, yes, outsourcing, represent today's versions of methodologies.

Outsourcing

Although Software Development '99 West isn't the type of show to focus on outsourcing—after all, the audience is made up of software developers who presumably are interested in doing their own development, not farming it out—it seems to be gaining in popularity. For example, at Stamford, CT-based Gartner Group Inc.'s combined Spring Symposium/ITxpo and Dataquest '99 Predicts held in San Diego, CA, in March, Paul Pastrone, vice president and worldwide director of software, systems, peripherals and Internet research at Gartner Group, observed that the combination of offshore outsourcing

and the proliferation of power-user products have conspired to move programming from the almost mystical status of high art to the commodity status it now has. Today, says Pastrone, business programmers are typically held in higher regard for their business knowledge rather than their coding skills.

Of course, custom consulting and contracting are as old as programming. But the pace certainly seems to have picked up. Not only are IT shops contracting out more "just-in-time" development (often a subset of a larger project), some organizations are outsourcing the entire IT function. Savvy developers need to be flexible enough to realize they could be outsourced out of a job. Savvy IT shops need to realize that, just as it may be more cost-effective to buy rather than build, it may also be smarter to outsource rather than integrate.

Part of this trend toward outsourcing can be attributed to the Internet ("The Internet changes everything," was the theme of Oracle's May iDevelop '99 conference held in Burlingame, CA), part of it to shorter cycle times and part of it toward restructuring that sees organizations shucking off all but their core competencies.

Oracle Chief Executive Officer, Larry Ellison, may have been one of the first to take the outsourcing trend to another level when he announced his vision for Oracle Business OnLine application hosting in 1998. As the company explains (see http://www.oracle.com/businessonline), "Oracle Business OnLine is an innovative service that delivers business applications and technology through a wire, all hosted on professionally managed servers. No need to procure expensive equipment. No need to attract and train hard-to-find IT staff. Just a simple service contract and a phone line." Does this sound like modern-day timesharing? Well, you know what they say, "Plus ça change, plus c'est la même chose."

Other firms are talking about software rental, sometimes referred to as application service hosting. The trend toward outsourcing and third-party application hosting doesn't necessarily mean that today's software developers risk losing their jobs but, combined

with the trend toward using more commercial off-the-shelf software, developers may find themselves either working for large systems integrators and service providers, such as IBM Corp.'s Global Services and Electronic Data Systems Corp., or simply providing for-hire labor as independent just-in-time contractors.

Masood Jabbar, president of Sun Microsystems Inc.'s Computer Systems Division, Palo Alto, CA, predicts that by 2002, it will all be about services. "I will not have to do my own application development," Jabbar says. Rather than roll his own enterprise resource planning (ERP) system (or rather, head a large team of 500 employees dedicated to implementing Sun's ERP system), Jabbar says he'll be able to call an ERP application provider—for example, Baan Co., Oracle, PeopleSoft Inc. or SAP AG—and say, "I need you to provision me with MRP [manufacturing resource planning], some financials, some human resources package." Ironically, another

part of Sun is already active in the application provisioning services sector, providing Web hosting services.

ASP, which refers to Active Server Pages in the Microsoft world, is increasingly being used to refer to the emerging space of "application service providers." It will be interesting to see how the market for application hosting/software rental pans out—and how it affects programmers in IT shops. Some will undoubtedly be redeployed to other projects, but many will likely be laid off.

Languages, Suites, Frameworks

Although Java and C++ were the hot languages at this year's show, they're not the only ones associated with object-oriented development. Visual Basic, while admittedly not a "pure" object-oriented programming language, has features built into it that allow millions of programmers to create ActiveX controls, which can be used by Windows programs that are Component Object Model (COM)-enabled. In fact, the

Microsoft Office 2000 suite extends this capability to Word and Excel users (see "Microsoft's Interoperability Efforts").

But as far as IT shops go, the trend seems to be to standardize on one (or more) software development suite. Leading examples are Microsoft's Visual Studio and IBM's VisualAge family. Visual Studio comprises Visual C++, Visual Basic, Visual FoxPro, Visual J++, Visual InterDev, Visual SourceSafe, the new Microsoft Data Engine (MSDE, essentially a stripped down version of Microsoft SQL Server) and various utilities such as Visual Component Manager and Microsoft Repository. Earlier this summer, Microsoft began promoting the fact that Visual Studio is integrated with Rational's RationalSuite product line (TestStudio, AnalystStudio and DevelopmentStudio). The latter includes UML-enabled Rational Rose, Rational Purify, Rational PureCoverage, Rational Quantify, Rational Requisite-Pro, Rational ClearQuest, Rational SoDA and Rational Unified Process.

Microsoft's Interoperability Efforts

At two Microsoft Corp. TechEd '99 sessions on Windows interoperability (you can view slides of the sessions at `http://www.teched99.com`), Aubrey Edwards, a member of Microsoft's Platform Integration Group, outlined the company's efforts to make Component Object Model (COM) available on non-Windows platforms. Today, you can obtain COM for Sun Microsystems Inc.'s Solaris operating system, COM interoperability products for Digital UNIX and OpenVMS from Compaq Computer Corp. and COM for certain other operating system environments from SAGA Software Inc. (formerly Software AG). Microsoft, however, has a richer port to Solaris under development and is also working to port COM to Hewlett-Packard Co.'s HP-UX.

Today, Microsoft Transaction Server, or MTS, (see Figure 1) can invoke XA-based resource managers (typically RDBMSs or print servers) for which there's an XA-compliant Open Database Connectivity (ODBC)/Object Linking and Embedding (OLE) database driver available (the Open Group's XA interface is the bidirectional interface between a transaction manager and resource managers). According to Edwards, MTS can also interoperate with Unisys' OpenTI and ClearPath/NX, BEA Systems Inc.'s BEA Tuxedo, Iona Technologies Inc.'s Orbix OTM and Compaq's Digital ACMS.

Edwards' final point in both presentations was that eXtensible Markup Language (XML) and Microsoft's BizTalk initiatives promise to make cross-platform interoperability easier. Perhaps that will reduce the need to port COM, DCOM and COM+ to non-Windows platforms—something that's happening today, but certainly not in a particularly timely fashion.—*kw*

Figure 1. Microsoft Distributed Transaction Architecture

Source: *Principles of Transaction Processing*, by Philip Bernstein and Eric Newcomer (published by Morgan Kaufman, 1997, ISBN 1-558-60415-4)

Microsoft's implementation has the same components as The Open Group's X/Open transaction model, but instead defines proprietary interfaces. Under the X/Open model, the transaction manager runs two-phase commit for all transactions on the network and communicates with local resource managers and remote transaction managers.

IBM's VisualAge family has even more members than Microsoft's Visual Studio, but unlike the latter, isn't available as a single suite. In addition to the popular VisualAge for Java and VisualAge C++, IBM also sells VisualAge for Embedded Systems, VisualAge Interspace, VisualAge COBOL, VisualAge 2000, VisualAge Developer Domain, VisualAge Pacbase, VisualAge PL/I, VisualAge RPG and VisualAge Smalltalk.

In addition to the language-specific VisualAge products, IBM also sells a version control product called VisualAge TeamConnection Enterprise Server, which is basically IBM's equivalent to the combination of Microsoft's SourceSafe and Repository tools, and a middleware product called the IBM Component Broker, which competes with Microsoft

Transaction Server (MTS). MTS is currently available as part of the Windows NT 4.0 Option Pack, but will be built into Windows 2000 as part of COM+.

It's worth noting that the most popular VisualAge tool is VisualAge for Java. Like IBM, Oracle too is strongly committed to Java, made clear by its recent release of jDeveloper suite. And all of these vendors are in the process of

Component Stew

This isn't a comprehensive list of components, but its length and variation in granularity should give you an idea about the challenges of implementing systems where objects can talk to one another.

Forms, reports — Large granularity objects that can be stored and re-invoked.

Component Object Model (COM) — A specification developed by Microsoft Corp. for building components that can be assembled into programs or used to add functionality to existing programs that run under Windows. COM components can be written in any number of languages, but are typically written in C.

COM object — COM objects refer to COM and can be small or large. They can be written in several programming languages, and they can perform any kind of processing. A program can call the object whenever it needs its services. Objects can be run remotely over the network in a distributed objects environment (DCOM). Depending on their usage, COM objects are also called ActiveX components, OCXs, OLE controls, or simply controls.

COM frameworks — Megalibraries that are available to developers who want to create COM objects. COM frameworks include Microsoft Foundation Classes (MFC) and the leaner Active Template Library (ATL).

Templates — Similar to class libraries, templates are typically supplied only as source code and not necessarily hierarchical. Rather than deriving from a class, you instantiate a class from a template. As with a macro, invoking a template causes it to expand (with appropriate parameter substitution) to code you have written.

Class factory, factory object — A class factory is a special type of object based on the factory pattern in which the code that actually creates the object is contained in its own binary file. COM servers provide class factories so that COM objects can be instantiated.

MTS object — An ActiveX control that runs under Microsoft Transaction Server (MTS).

Applet — A program, usually written in Java, that can be downloaded from the Web and run on a client system. A Java applet is a Java class that is loaded and run by a running application such as a Java Virtual Machine (JVM).

Servlet — A Java program that runs off a server. Conceptually similar to CGI scripts, which control interaction between Web servers and resources on host computers.

Plug-in — A program that "plugs into" another, usually larger, application to provide added functionality. Plug-in components are often associated with the Netscape Communications Corp. Web browser, and are often available free as "readers" for proprietary formats.

Dynamic Link Library (DLL) — A feature of Windows and OS/2 operating systems that allows executable routines to be stored as files with .dll extensions and loaded only when needed.

Stored procedure — A series of Structured Query Language (SQL) statements that are precompiled and stored in the database server.

Active Server Pages (ASP), JavaServer Pages (JSP) — Text files that contain HTML tags, text and commands written in a scripting language and stored on a server. ASP and JSP can contain references to supported components and can be used to control a client's interaction with a Web page.

JavaBeans — A Java component architecture defined by Sun Microsystems Inc., where JavaBeans are reusable application components similar to Microsoft ActiveX controls. JavaBeans consist of one or more Java classes (which themselves can consist of one or more Java objects) and Java messages (instructions). JavaBeans can interrogate each other and be interrogated, a feature known as "introspection" or "reflection."

Enterprise JavaBeans (EJB) — Sun's component architecture for distributed applications. A specification (not a code library) that describes what amounts to JavaBeans for servers.

Class library — A set of ready-made software routines (class definitions) that programmers use for writing object-oriented programs. These class definitions also include their inheritance characteristics, if applicable.

Java Foundation Classes (JFC) — A class library from Sun that provides an application framework and GUI routines for Java programmers. Sun, Netscape, IBM Corp. and others contributed to JFC, which combines Sun's Abstract Windowing Toolkit (AWT) and Netscape's Internet Foundation Classes (IFC).

Agent, bot, intelligent agent, spider — Components (programs) that perform background tasks such as searching for information over the Internet or otherwise mediating communication between clients and servers.

Common Object Request Broker Architecture (CORBA) — A distributed object mechanism by which objects written in different languages and executed on different platforms can all communicate. The CORBA standard is published by the Object Management Group, a consortium of industry vendors.

Object Request Broker (ORB) — The basic "middleware" entity specified by CORBA. Several commercial ORBs exist, and an increasing number of application vendors are building applications that can "talk" to ORBs.

Interface Definition Language (IDL) — CORBA's language for describing an object (service) in CORBA. Microsoft's IDL for MTS is called MIDL. CORBA has various IDL compilers for different languages.

weaving, or welding, eXtensible Markup Language (XML) into their suites and/or relational database management systems (RDBMs).

What about fourth-generation languages (4GLs)? Are they dead? Do they have any role in an object-oriented world? No, 4GLs aren't dead. In fact, although most 4GL development targets existing customer bases, some 4GLs are attracting new customers.

One advantage to object-oriented programming is that it makes it relatively easy to develop and roll out projects incrementally, and that's turning out to be a lifesaver for many organizations.

4GLs, you may remember, were introduced in the late '70s as languages that were supposed to increase the productivity of third-generation COBOL and FORTRAN programmers by a factor of 10. Information Builders Inc.'s Focus, Progress Software Corp.'s Progress 4GL and Cognos Inc.'s PowerHouse 4GL are all examples of 4GLs that continue to have significant numbers of loyal customers, many of whom are migrating to more object-oriented versions of the vendors' products (WebFocus, Apptivity and PowerHouse Web, respectively). Other more PC-oriented development suites that openly tout themselves as 4GLs include Omnis Studio from Omnis Software Ltd. and Clarion from TopSpeed Corp.

And what about client/server? Is it dead? According to Oracle's Ellison it is. Client/server, largely associated with two-tier architectures ("fat" front-end clients, which Microsoft began referring to as "rich" clients earlier this year, and back-end database servers), simply doesn't have a place in the Internet economy of browser-enabled clients

that interact with any number of Web, application and database servers. But, despite Ellison's druthers, not everyone is ready to dump their inventory of battle-tested client/server applications, nor are they anxious to run their entire company over the Internet. No, client/server isn't dead, but it will probably fade away—mainly because of the high cost of supporting Windows fat clients.

ERP, E-Business Drive New Development

Most new object development is being done to create and evolve e-business sites, but Personal Digital Assistants (PDAs), such as 3Com Corp.'s PalmPilot, Windows CE devices and a host of special-purpose embedded devices, are also beginning to represent significant new development. In his Software Development West '99 keynote address, 1978 Nobel Prize Laureate in Physics Arno A. Penzias, predicted that we're moving beyond the keyboard and mouse to computers that have "eyes and ears." Clearly, the availability of pervasive con-

tinuous speech recognition technology and the vastly increased use of real-time video camera feeds and global positioning satellite technology will also drive software development.

One advantage to object-oriented programming is that it makes it relatively easy to develop and roll out projects incrementally, and that's turning out to be a lifesaver for many organizations. This means, for example, that development teams don't have to delay the

deployment of an e-business site until they've got their entire solution developed and tested. Instead, they can test the waters with a business-to-business extranet before completing integration with their ERP and customer relationship management (CRM) systems and before launching a retail electronic commerce site. Object-oriented development also eases the pain of the periodic wholesale facelifts most organizations have to do to support new browser functionality.

Of course, a good deal of Internet development activity is being outsourced. According to *Business Week*'s "Information Technology Annual Report," released June 21, IBM Chief Executive Officer Lou Gerstner told Wall Street analysts in May that market researchers predict 60% of the money spent on e-business industrywide will be for services.

One benefit of the frenetic pace of global e-business is that standards are being created and adopted at a heretofore unthinkable pace. After all, objects and components from one site's application need to be able to "talk" to other

Table 1. Object Models: 1999 Report Card

	CORBA	COM/COM+	EJB
MATURITY	A	B	C
AVAILABILITY	A	D	B
IDE SUPPORT	D	A	A
ISV ENTHUSIASM	C	A	A
CHAMPION(S)	C	B	A
COMPONENT LIBRARY	D	B	B
LANGUAGE(S)	A	A	D
FIVE-YEAR PLAN	D	B	B
OVERALL	B–	B	B+

Source: Gartner Group Inc.

sites' application components. Which brings us to ActiveX objects, Enterprise JavaBeans (EJB) and the state of the COM/CORBA "wars."

COM, CORBA and EJB

It's no secret that Java and the Web legitimized–and accelerated–the use of components. The idea behind components is that whether you build and test your own, have someone else build them for you, or you simply buy them "off the shelf," they can be used over and over again.

Components have been around in

one guise or another for some time (see "Component Stew," Page 54), but Microsoft's Visual Basic eXtensions (VBX) were the first widely used components. VBX were originally limited to use in 16-bit Visual Basic programs–meaning that you could only use them in a Windows environment. In fact, VBX were really just 16-bit Windows Dynamic Link Libraries (DLLs), so for all practical purposes, you had to be a C/C++ programmer to create them. But any one of the millions of Visual Basic developers could use them. Today, even Office 2000 users can create the current generation of Microsoft components called ActiveX controls.

ActiveX controls, which are COM objects, know something about themselves and can be queried in a standard, well-defined fashion. The problem with ActiveX controls and their underlying COM technology is that they're basically limited to the Microsoft Windows environment. Hence, the appeal of Enter-

Companies Mentioned in this Article

BEA Systems Inc.
2315 N. First St.
San Jose, CA 95131
http://www.beasys.com
Circle 140

Cognos Inc.
P.O. Box 9707
Ottawa, Ontario
Canada K1G 4K9
http://www.cognos.com
Circle 141

Compaq Computer Corp.
P.O. Box 692000
Houston, TX 77269
http://www.compaq.com
Circle 142

Electronic Data Systems Corp.
5400 Legacy Drive
Plano, TX 75024
http://www.eds.com
Circle 143

IBM Corp.
Contact local sales office
http://www.ibm.com

Information Builders Inc.
2 Penn Plaza
New York, NY 10121
http://www.ibi.com
Circle 144

Inprise Corp.
100 Enterprise Way
Scotts Valley, CA 95066
http://www.inprise.com
Circle 145

Iona Technologies Inc
60 Aberdeen Ave.
Cambridge, MA 02138
http://www.iona.com
Circle 146

Microsoft Corp.
1 Microsoft Way
Redmond, WA 98052
http://www.microsoft.com
Circle 147

Omnis Software Ltd.
Salamander Quay W.
Park Lane, Harefield
Middlesex, U.K. UB9 6NZ
http://www.omnis-software.com
Circle 148

Oracle Corp.
500 Oracle Pkwy.
Redwood Shores, CA 94065
http://www.oracle.com
Circle 149

Progress Software Corp.
14 Oak Park
Bedford, MA 01730
http://www.progress.com
Circle 150

Rational Software Corp.
18880 Homestead Road
Cupertino, CA 95014
http://www.rational.com
Circle 151

SAGA Software Inc.
11190 Sunrise Valley Drive
Reston, VA 20191
http://www.sagasoftware.com
Circle 152

Sun Microsystems Inc.
901 San Antonio Road
Palo Alto, CA 94303
http://www.sun.com
Circle 153

TopSpeed Corp.
150 E. Sample Road
Pompano Beach, FL 33064
http://www.clarion.com
Circle 154

Visual Edge Ltd.
3950 Côte Vertu
St-Laurent, Québec
Canada H4R 1V4
http://www.visualedge.com
Circle 155

prise JavaBeans (EJB) and OMG's cross-platform Common Object Request Broker Architecture (CORBA). Table 1 illustrates Gartner Group analyst John Enck's assessment of the fitness of the competing distributed object models for high-end enterprise applications (as opposed to tactical or mid-size applications).

EJB, whose specification has been developed by a consortium of companies led by Sun, is an API that extends the JavaBeans component model to cross-platform, server-side applications. Like COM objects, EJB, which come in two major flavors—session and entity beans—know something about themselves and are capable of introspection.

The problem with ActiveX controls and their underlying COM technology is that they're basically limited to the Microsoft Windows environment.

CORBA, on the other hand, is a rich but complex, decade-old component-based middleware architecture that supports multiple languages. In his presentation on Microsoft's distributed Internet architecutre (DNA) at the Gartner Group's Windows NT conference held in Indian Wells, CA, in May, Enck predicted that "over time, the EJB and CORBA models will merge, as each model addresses the weaknesses in the other." EJB, says Enck, will "lose the Java" and evolve into EB, or Enterprise Beans, ultimately replacing CORBA. Enck says IBM is best positioned to champion EJB/EB enterprise usage.

Fortunately, Microsoft's June investment in Inprise Corp. (formerly Borland) gives the company access to Inprise's CORBA-compliant Visigenic ORB (object request broker) and COM/CORBA bridge technology. Other vendors, including BEA Systems Inc., Iona Technologies Inc. and Visual Edge Ltd., already offer COM/CORBA bridges. Assuming that Microsoft uses the

investment to build a version of COM+ (the next generation of COM built into Windows 2000) that can talk to CORBA object brokers and EJB, the COM/CORBA wars may end in peace.

Today's Challenge: Fast, Good and Flexible

The days of multiyear systems projects are all but gone. Now, more than ever before, IT is expected to deliver solutions in weeks or, at most, months. Modular, component-based development helps IT deliver functionality, often incrementally.

Some jaundiced developers see this "incremental" development paradigm as yet another manifestation of Pareto's 80/20 rule. (Alfred Pareto was a 19th century Italian scientist who created the distribution analysis that bears his name and which we often call the 80/20 rule, or ABC classification. His original survey revealed that 80% of the wealth of his country was in hands of 20% of the population.) Is it ethical to deliver software that only offers an estimated 80% functionality? Or is "good enough," "quick and dirty" software development destined to be the norm?

There are other challenges associated with trying to deliver software today: testing complex assemblies that may consist of hundreds or even thousands of components—and living with the complex dependencies that exist among them. Although testing suites are evolving to address these challenges, ISVs and corporate developers alike may lack the discipline to perform the necessary testing, especially in the face of the competitive pressures of the Internet. (Note: Software Development '99 East will be held at the Washington Convention Center in Washington, D.C., November 8-12.) ➝

Karen Watterson is an independent consultant and author based in San Diego, CA, who specializes in database and data warehouse design. She recently coauthored *10 Projects You Can Do with Microsoft SQL Server 7* (published by John Wiley & Sons, 1999, ISBN 0-471-32751-4). Email: karen_watterson@ msn.com.

The Time of SANs

Despite the recent hoopla, technologies necessary for wide SAN deployment are only just beginning to fall into place.

This is the year storage area networks (SANs) will become the architectural standard for enterprise-wide network storage. The SAN, a high-speed, broad bandwidth I/O channel that connects to the back end of local area network (LAN) servers, promises to improve overall corporate network performance as it takes its place alongside an organization's LAN and WAN. SAN nodes are a mixture of storage devices (for example, RAID and automated tapes) capable of communicating directly with one another over a fast, fault-tolerant storage pipeline (the SAN) with multiple hosts.

A typical SAN includes a high-throughput interface (Fibre Channel being the most common), interconnects (switches, gateways and hubs) and a protocol (SCSI, IP or ESCON). The SAN is a shared storage repository attached to multiple servers via an independent network, which could potentially remove all storage functions from the LAN and the WAN.

Why SANs?

SANs are needed because many businesses and organizations now rely on UNIX and Windows NT distributed platforms to run crucial enterprise applications. This shift from the data center to the distributed network is possible because software and processor performance have both increased manyfold over the

BY RON LEVINE

HOURGLASS PHOTO BY PHOTODISC INC./MONTAGE BY BRAD DILLMAN

past decade. However, storage devices and storage interfaces–despite some improvements–have lagged far behind in terms of performance. This disparity in functionality often results in network bottlenecks at sites where high storage access rates or large volume file transfers are the norm.

But today, storage functionality can be scaled to be on a par with processor and software performance by integrating a SAN infrastructure into an existing network.

A SAN basically extends the any-to-any connectivity of LAN architectures to storage resources. Any data on the network, in any location, can be made accessible through multiple paths to any node, application and user on the network. The ubiquitous data access provided by this new paradigm completely changes the economics of availability and scalability, and offers other manageability advantages that leverage the centralization of storage management.

Once connected using a SAN topology, storage devices are in effect uncoupled from the restrictive dedicated SCSI bus server-to-individual storage unit connection. By creating multiserver-accessible and shared-network storage devices, all storage and file access tasks are freed from the operational limitations of an individual server and the general traffic constraints of a LAN or a WAN. As a result, operational improvements to a storage subsystem's performance and a network's overall performance can be realized. In addition, because a SAN externalizes storage outside the server, the stored data is made available to multiple hosts across multiple platforms without imposing extra overhead on the network.

Current SAN Benefits – Available Today
• Increased bandwidth
• Increased speed (throughput)
• Improved overall performance of storage subsystem
• Scalable storage capacity
• Improved availability
• Enhanced fault-tolerant storage operations
• Elimination of distance restrictions present with SCSI connections
• Removal of LAN traffic bottlenecks caused by storage tasks
• Data sharing among multiple servers and applications

Future Benefits – On the Horizon
• Centralized storage administration
• Dynamic redistribution of storage resources
• Streamlined data management processes
• Data-centric architectures

Stage One SAN vs. SCSI

BENEFIT	SCSI	FIBRE CHANNEL
Higher bandwidth	80 MB/s (max.)	100 MB/s
Channel distance (max.)	25 meters	10 kilometers
Number of devices	15 nodes/channel	126 nodes/loop, 16 million devices/fabric

A SAN basically extends the any-to-any connectivity of LAN architectures to storage resources. Any data on the network, in any location, can be made accessible through multiple paths to any node, application and user on the network.

Incorporating a SAN into the enterprise network environment enhances LAN/WAN throughput and reduces traffic jams; the communications network is relieved of the bulky overhead associated with file access, file retrieval, data storage and data backup functions–tasks that generally hog network bandwidth. Users see immediate storage access and storage availability improvements because file read/writes, backup/

restore, archiving/retrieval, data migration and data/device sharing are more efficiently handled by a network optimized for the high-throughput, large-packet data transfer demands of storage tasks.

Behind the scenes, the storage subsystem's reliability is also improved because of the SAN's inherent fault-tolerant operation–with a SAN, data is accessible via alternate data paths (dual ports, for example). Redundant access can easily be built in. Administration and management of a network's storage components are simplified because a SAN creates an environment where relatively easy scaling, servicing and management of storage devices and their data is possible because of centralization. This centralization significantly improves backup and restore capabilities, and makes it easier to scale and achieve cost-effective disaster protection configurations.

SAN Benefits

Any site deploying business-critical applications over distributed networks must treat the associated data as a vital company asset and should consider a SAN implementation. The uninterrupted availability and reliability of this data can be guaranteed by integrating a SAN into an organization's existing computing infrastructure. Because SANs provide fault-tolerant storage operations, they ensure that stored data remains intact and available at all times–even with the loss of a major network component, such as a server.

Besides data storage, file access and functional reliability protections, there are other reasons to install a SAN. Scott

Enterprise-Class SAN Solutions
Tested, Optimized, Certified. . . Guaranteed

The pundits will tell you that the promise of SAN technology remains in the future. Maybe from other vendors. At DataDirect Networks we've already integrated some of the world's largest, fastest, and most scalable SAN's. Our technology integrates all four of the necessary components to enable a cross-platform, scalable, enterprise-class SAN that unleashes the power and creativity held captive in your organization. Our SAN file systems, Network Performance Software, Fibre Channel RAID storage, and Fibre Channel Network designs deliver your information - when you want it, where you want it.

When you're serious about the benefits of incorporating a SAN, there's only one place to turn. Call DataDirect Networks today to request your complimentary copy of "The DataDirect Networks' SAN Solution Advantage".

Information when you want it, where you want it

dataDIRECT NETWORKS

ales@datadirectnet.com • 800.322.4744 • 818.700.7600 • fax 818.700.7601 • www.datadirectnet.com

Circle No. 18

Robinson, vice president of engineering at Minneapolis, MN-based Datalink Corp., points out that even a stage one (Fibre Channel bus) SAN implementation can improve overall network operations. The resulting increased bandwidth capabilities greatly enhance I/O and network performance at any site where high-end throughput applications–such as imaging, video and CAD/CAM–are dragging down overall network responsiveness. Likewise, a SAN's Fibre Channel interface improves the response time of any network running high-end I/O applications like OLTP and databases. Robinson says additional benefits, such as zoning and high-availability clusters, will further improve overall network performance as auxiliary SAN components become available later this year. And, with a SAN, it's relatively painless to add extra storage capacity, independent of any specific server or operating system platform on the network.

The ability to isolate storage functions from one another, to easily scale storage capacities and to add/remove storage units is now possible through the implementation of Fibre Channel switches and hubs.

In general, SANs make high-availability storage subsystems more cost-effective for a wider variety of application environments. And the hot-pluggable nature of SAN resources enables scalability, while preserving ubiquitous data access. With storage tasks assigned to a dedicated network, more efficient management of optimization, reconfiguration and backup/restore data becomes possible.

SAN Time Line

While 1999 is shaping up to be the year corporate America embraces SAN technology, the full potential of SANs still cannot be fully realized. All of the pieces (hardware and software) necessary for a true SAN environment are not yet available, but they're beginning to appear.

According to Datalink's Robinson, SANs are being rolled out in three stages:

STAGE ONE
1998-1999 Stage one SAN installations exist now and are delivering some immediate and measurable operational benefits. Stage one is designed to speed up the performance of common storage device-to-network data transfers. By connecting RAID and other storage devices over Fibre Channel buses, higher bandwidth and increased speed I/O transfers can occur over longer distances than with a standard SCSI interface. The RAID and Fibre Channel combination improves storage subsystem reliability through fault-tolerant disk operations and pipeline (channel) connections.

The ability to isolate storage functions from one another, to easily scale storage capacities (independent of server func-

tionality) and to add/remove storage units (without affecting network operations) is now possible through the implementation of Fibre Channel switches and hubs.

At sites where Fibre Channel-based SANs are already in operation, SAN-based data sharing among various applications is a reality. And other benefits include reduced storage administration time owing to the centralized data management environment and the ease with which backup procedures can be implemented. Stage one SAN technology is alleviating storage-related network bottlenecks and easing management of the network's storage component.

STAGE TWO
1999-2000 Stage two SAN functions are beginning to be introduced. Some components are available, while others will be out later this year or early next year. SAN-specific capabilities like zoning, Logical Unit Number (LUN) masking and high-availability clusters will soon be implemented on Fibre Channel fabrics.

Zoning is the process of assigning a storage resource (for example, a disk) to one server, while another resource (say, a tape) is committed to another server. Fibre Channel switches permit storage devices to be connected to dedicated storage zones, providing the benefits of zone fault isolation–the capability to dynamically add/remove devices to a server's storage pool while maintaining a full 100-MB/s bandwidth.

Software will soon be made available that will take advantage of the dynamic rezoning of storage devices. For example, with zoning, a tape device can be temporarily included into the disk device's zone during a specific backup window for direct-attached backup. Then, it can be reallocated to a different zone for remote vaulting or archiving. To accomplish this task, the backup software must be able to recognize the dynamic appearance of a shared tape resource across multiple servers. Software to support this concept is expected to begin shipping by the end of the year.

Logical Unit Number (LUN) masking subdivides a storage device (such as RAID) into individual sections, creating a shared resource across multiple network servers. By appropriating individual storage space within a device, file sharing

The SAN Time Line

Early Virtualization
Logical Volumes
PPRC
Outboard Backup
Loop Extension
LUN Masking
Zoning
Tape Sharing
Switched Fabric
Loop
Point-to-Point

1998 1999 2000 2001 2002 2003

Gartner Group

among servers and workgroups can be supported on a selected basis. These assigned storage areas can also be earmarked for specified applications or file types and information-specific access management.

Also making their debut in 1999 are high-availability and high-scalability SAN clusters. As sites replace older SCSI I/O interfaces with Fibre Channel pipelines and SAN support software explodes on the scene, sturdy and robust cluster configurations will be installed. SCSI-based clusters aren't capable of being enhanced with the performance and scalability features of Fibre Channel-based clusters. These new, more powerful clusters (built on Fibre channel technology) will enable the dynamic reallocation of storage resources, furnishing fault isolation on the SAN fabric. Also, Fibre Channel and SAN cabling options allow the physical separation of redundant components along the SAN. These new clusters will also permit the adding/removing of storage devices without affecting network operations.

STAGE THREE
2000+

Stage three is still several years away, Robinson says. "I see stage three SAN development as bringing a major shift in the way IT professionals think about information systems. I see network administrators placing information management needs ahead of LAN-management needs when configuring their IS operations," he says.

Software will be the key to effectively integrating storage hardware technologies and delivering on the promise of SANs. SAN benefits will be achieved with the blending of an array of software applications for functions such as enterprisewide storage management, high-availability systems, network backup and archiving, disaster protection and Hierarchical Storage Management (HSM) with diversified storage hardware devices like RAID, tape, optical disk and DVD.

Based on real-world discussions with corporate clients, Robinson says the move toward information-centric organizations will fuel the rapid growth in information sharing across multiple platforms and applications. The solid integration and tight coupling of numerous software functions, such as file systems, volume managers, HSM and archiving, will further improve and simplify enterprisewide storage administration and management, while encouraging enterprisewide data sharing.

Eric Burgener, group manager of SAN products at Veritas Software Corp., a Mountain View, CA-based vendor of heterogeneous storage management tools, agrees with Robinson's assessment. Veritas also adheres to a three-step general industry rollout of SANs. Its time line looks like this:

1999 – Storage appliances; local SANs; homogeneous device sharing.

2000 – Hardware API standards; wide area SANs; homogeneous data sharing.

2001 – Operating system API standards; global SANs; heterogeneous data sharing.

In the first step, vendors are providing the hardware and software components necessary to support pluggable storage resources, called appliances, in local SANs. In a local SAN, all nodes are within the same building, but not necessarily on the same network subnets. Homogeneous device sharing allows separate nodes running the same operating system to access different data sets on the same physical storage device.

Burgener envisions the second step as the development of hardware APIs, which will allow the proliferation of cost-effective, yet high-performance, commodity storage appliances targeted at specific application environments. Movement to the wide area through the use of Fibre Channel-attached devices and replication technologies will allow SANs to be harnessed for disaster recovery and other multisite configurations. Cost-effective and easy-to-implement products will support off-the-shelf applications with shared concurrent access to storage resources from nodes running the same operating system environment. This is known as homogeneous data sharing.

Third step projections call for volume access protocols implemented through standard APIs at the kernel layer of an operating system, which will support heterogeneous data sharing. In this model, different nodes running different operating systems will have shared concurrent access to the same underlying data set. The increasing deployment of high-speed networks, combined with enhancements in replication technologies, will allow the implementation of high-performance global SANs.

These developments will pave the way for easier and more ubiquitous access to storage resources, allowing customers to more fully leverage their cross-platform investments against a highly available, scalable and centrally managed common data store.

Datalink's Robinson believes that what lies ahead is a world of multinode heterogeneous SAN clusters, which will enable the rapid growth of shared storage application functions in a high-availability IT environment. �థ

Ron Levine is a technology writer with Coast Writing, an independent firm that specializes in computer application articles. He can be reached at ron@coastwriting.com.

Applixware:
A Fine Alternative

by IAN DARWIN

Tired of all the bugs in your current office software?
Then maybe the platform-independent Applixware
office suite is just what you're looking for.

It had been a stressful week. A big project was due to one client, while another client was complaining about changes to their Web site. The first client had insisted I write their documentation in Microsoft Corp. Word 97 under Windows 95. They'd even given me a shrink-wrapped copy of Office 97. But by the time it had bombed out on me for the second time in one day, a little voice in my head was saying, "This software is worth exactly what you paid for it."

Each time the software died, I swore it'd be the last, as I went painstakingly back through the file to find out if my changes had been saved (in Word's favor, most of my changes were saved; but what do you get for the time spent rechecking all those changes?). It died again. And again. I yelled, "Strike three!" (In truth, I yelled something considerably less printable.) This time, instead of restarting Word, I started my browser and went straight to Applix Inc.'s Web site and ordered a copy of the Applixware office suite. It promises platform independence and the ability to import and export Office 97 files, along with many other formats.

Applixware runs on various flavors of UNIX and Windows 95/NT. The Windows 95 version was not evaluated here because then I'd have to admit to still having a Microsoft operating system on one of my systems, which I'd rather not do. Applixware is available for the following platforms: Solaris 2.5+ (SPARC or Intel), Linux (Intel, PowerPC, DEC/ Compaq Alpha), HP-UX 10.2, AIX 4.2, Digital UNIX 4.0, IRIX 6.2 and Windows 95 (Intel) or NT 4. Some of the internal code is written in Java, which simplifies the porting of code to new platforms.

To make a long story short, Applixware generally fulfilled the company's claims and met (or exceeded) most of my expectations. And yes, I have since removed or excluded Word from all of my own systems, plus a few others I control.

Installation

I tested Applixware on two platforms. If you remember my review of OpenBSD ("A Good Web Server Platform at the Right Price," June 1998, Page 81, http://sw.expert.com/ R/WS4.JUN.98.pdf), you'll know it's my chosen freeware UNIX. Indeed, the Linux port of Applixware runs perfectly in Linux emulation mode on a Pentium under OpenBSD. I also tested it on an old Sun Microsystems Inc. SPARCstation IPX running Solaris 2.5.1. Apart from the installation, I was unable to ascertain any difference.

Figure 1 shows both versions' toolbar windows and "about version" windows running side by side. If you can spot any differences–other than one has a license number blurred out for security reasons and the other does not–consult an optician.

Installation, though, was slightly different. Linux installation has a setup

Figure 1. Toolbar Window Comparison

shell script front end included in the CD-ROM, which decides between "libc5" or "libc6" versions, and runs the appropriate compiled executable. For the SPARC version, you just `cd` to `setup/applix` and run the `install` executable.

As usual, when people try to automate UNIX installations, simple installs are simple and nonsimple cases are not. Both versions failed to judge free space adequately when installing into an NFS-mounted partition. Not only that, they responded differently. The Linux version gave a `df`-like listing of partitions with free space; none of which matched the reported free space (see Listing 1). The Solaris version simply gave up, reporting numbers like 24,117 KB available, 27,398 KB required, with no indication of which file system was in trouble. In the end, I gave up on NFS installation by rearranging the file systems so that each machine had enough space (170 MB minimum!) to install on a local disk. It seems no matter how much disk space you have, installing new software turns into an excuse to delete old, unused software.

Once I made the required space available locally, the installation went smoothly, even with the CD-ROM NFS-mounted on one of the systems.

Incidentally, both CD-ROMs have a binary in the root of the CD; so to try it out, just `cd` into the CD and type `./applix&`.

The Solaris version can be licensed on either a node-locked or floating-license server; the latter uses a version of the familiar FlexLM license manager. The documentation includes a how-to-use book called *Make It Happen*. The Solaris version also includes a *System Administration Guide* that covers network installation and licensing. Both manuals seem adequate. (One amusing note: An indication that this product has been around for a while, Chapter 3, Page 20 of the administration guide refers to an email address `uunet!mv.us.adobe.com!ps_file_server`. Uhmm. It's been a while since Uunet was the center of the email UNIXverse and everybody had to do manual routing of email via the "!" path mechanism.) There are also several online books you can install.

Who Needs Microsoft?

Applixware is a comprehensive office suite that includes a word processing program, spreadsheet tool, graphics and slide show program, HTML editor, mail reader/composer and extensive customization. In addition, it can import from or export to many other formats (see Table 1). As you'll see, it can do pretty much everything that Microsoft Office does, without you having to resort to Windows emulation or running two operating systems.

Also, unlike most of its competitors, Applix native formats are ASCII-text-based, not binary. So it is somewhat closer to an open format, in the sense that you can `grep` through it with some success. I've even made tiny changes to documents using `vi`, something you would never think of doing to, say, a Word `.doc` file. Let's look at each of the major Applixware programs in turn, and make some comparisons with competitors.

Words

Since the Microsoft lawyers started grabbing up all the nouns on the planet, it has become increasingly difficult to find (nonlitigiously) a name for a new product. Applix chose "Words" for its word processing program. Applix Words is a word processor on a par with Word or Corel Corp. WordPerfect.

Unlike Adobe Systems Inc. FrameMaker, however, Applix Words has only one level of named styles, which contains both character and paragraph formatting. FrameMaker, by contrast, has separate, named styles for character and paragraph properties, which leads to

---

Listing 1. nfs-install-fail.txt

```
There is not enough disk space available in /home/applix.
175 MB are required, but only 14 MB are available.

Here is a list of local filesystems and the currently
available disk space:
=============================================================================
/                    (/dev/dsk/c0t3d0s0 ):              39722 blocks    21363 files
/usr                 (/dev/dsk/c0t3d0s6 ):              15260 blocks    82591 files
/proc                (/proc            ):                  0 blocks      436 files
/dev/fd              (fd               ):                  0 blocks        0 files
/var                 (/dev/dsk/c0t3d0s3 ):             101774 blocks    42660 files
/tmp                 (swap             ):              44056 blocks     5786 files
/opt                 (/dev/dsk/c0t3d0s4 ):             189570 blocks   303270 files
/var/mail            (darian:/var/mail  ):              21536 blocks    11121 files
/usr/src             (darian:/usr/src   ):             618432 blocks    97455 files
/usr/ports/distfiles (darian:/usr/ports/distfiles):     150672 blocks    90624 files
/stage/applix        (darian:/stage/applix):           228768 blocks   225178 files
/home/ian            (darian:/home/ian  ):             295040 blocks   108844 files
/cd0                 (darian:/cd0       ):                  0 blocks        0 files
/home/applix         (darian:/home/applix):            295040 blocks   108844 files
=============================================================================
```

Product Review

much better document structuring. If, for example, you have two styles named "code" and "literal" and you need to change all of code from Courier to LucidaSansTypewriter without also changing literal, you can do it with only a few keystrokes in FrameMaker. It's a manual, find-every-occurrence-and-decide operation with most other word processors.

The Applix Words spell check feature was uneventful. As with most such tools, its vocabulary is limited. Once you train it—which is as easy as clicking on "Add Word" in the spell check dialog box—it gets along fine, and does its job well.

My one area of concern is Words' import filter. Documents from Word-Perfect, Word and FrameMaker can be imported, although FrameMaker files must be saved in an interchange format known as Maker Interchange Format, or MIF. For simple documents, these imports worked fine. I imported an invoice form from an old DOS Word-Perfect file and it came out 99% correct (the only fault was that a couple of rules in one table were converted from double to single lines).

However, when I imported more complex documents from Word and FrameMaker, things didn't go so well. Because Applix Words doesn't have named page formats, it tends to drop parts of FrameMaker's "master pages" at the start of a document. Also, some graphics simply disappeared, or got dumped into the front of a file, when imported from Word. Both problems were fairly harmless, resulting in extra pages to delete, but even then, some of the paragraphs didn't import correctly.

Spreadsheets

Applix Spreadsheets, the spreadsheet processor, is a powerful spreadsheet program, comparable to Microsoft Excel. Applix Spreadsheets can import/export a variety of formats. I imported a simple Excel file and had no problems. When importing a more complex file (an expense report), I found that a graphics object wasn't aligned over the correct cells, but everything else came through OK. Applix Speadsheets has the features you'd expect in a spreadsheet program, including macros, built-in functions and charting capabilities, plus it has real-time spreadsheet capabilities that I didn't exercise.

Graphics and Presents

Graphics, Applixware's illustration/paint tool, and Presents, its slide show tool, are lumped together into a single program. Based on the file you wish to open, Applixware decides whether you are working on a single image or a slide show.

Presents is a standard slide show program. I've toyed with other slide show tools, including Microsoft PowerPoint, and this tool seems to have most of the same features as the others. Notably, its import filter worked very well for the slide shows I imported from Power-Point (see Figure 2). While the Spreadsheet import dialog isn't smart enough to select the file type when you click on a file, Presents is.

Figure 2. Presents Import Filter

The Rest of the Suite

Other programs that are part of the Applixware office suite, which I did not examine, include an email client (and OpenMail, which is X.400-compliant), Data (which gives access to relational databases), a directory browser, an ELF (Applixware's extension scripting language) macro editor and an HTML editor.

One powerful feature of Applixware is its general object-linking technology; this is certainly comparable to Microsoft Office's use of Distributed COM and/or ActiveX. For example, you can embed a spreadsheet in a presentation, or charts in a slide show.

Words also features hyperlinks and cross-reference features. This is quite general. If you import an HTML document into Words and click on what was an <A> tag in HTML, you can usually see the resulting file. For example, a link to a .gif file is opened in a separate Graphics window, and a link to another HTML file is opened

Table 1. File Formats Supported

PROGRAM	IMPORT	EXPORT
WORDS	ASCII, DCA, FrameMaker 4/5 (MIF only), HTML, Interleaf 4 (ASCII only), Word (up to 95/97), OfficeWriter, RTF, WordPerfect 4/5/6/7.	Applix 3, ASCII (various), FrameMaker MIF 4/5, Interleaf 4/5 (ASCII), Word (generally via RTF), WordPerfect 5/6.
SPREADSHEETS	WKS, XLS, SYLK, DIF/XDIF, CSV, ASCII (various, including import wizard).	XLS 3/4/5, WK1/WK3, Applix 3, CSV, DIF, SYLK.
GRAPHICS/PRESENTS	Applix, CGM, DXF, EMF, EPSI, FAX, GEM, GIF, HPGL, ILBM, JPEG, MacPaint, Windows BMP, PBM/PGM/PPM, PCX, PICT, PNG, PPT, Raw, SGI, SunRaster, TGA, TIFF, WMF.	Applix, CGM, EPSI, FAX, GIF87/89, HPGL, JPEG, Windows BMP, PBM/PGM/PPM, PowerPoint 97, SGI, SunRaster, TIFF lsb/msb, WMF, XBM (not XPM), XWD.

in…you guessed it, a separate Words window. So it's not a regular browser, but it does a good job of following and displaying links.

The installation program offers a number of optional extras. If you have a ton of disk space, install it all—you can always delete parts of it afterward. These include international dictionaries (spelling/hyphenation), hypertext help, online books describing the programs and programmability in more detail and clip art for Graphics and Presents.

In addition, there are quite a few extra fonts provided. The Solaris CD includes the fonts as BDF, SNF and even OpenWindows "FB." Reading the *System Administration Guide,* it seems difficult to add fonts and support additional printer features owing to quite a bit of editing. There is a font manager, but I didn't install it because I have plenty of fonts already.

Programming Applix

Applixware is highly customizable. There is an extension scripting language called ELF. It is lamentable that it has its own scripting language instead of using one of the "obvious" choices, such as Tcl or Perl. However, Applixware has been under continuous development since at least 1990, and what is obvious today may not have been obvious a decade ago. In the spirit of today's open-source movement, the ELF language, called SHELF, is being released as a scripting language in its own right; you can check this out at the Applixware Open Source Site (http://www.applixware.org).

In addition to ELF, there is a product called Builder, which allows you to build a complete application using ELF. Builder includes a GUI builder, which allows you to create the entire user interface visually, make connections and so on. So far, it is similar to, say, Sun's Open Look-based Developer's Guide or CDE Builder. But Applix Builder also allows you to write the actual application as a series of ELF calls. It seems, in fact, to have all of the development power of a tool like Microsoft Visual Basic. While I didn't write any full applications in it, Figure 3 is a snapshot of building a GUI.

ELF is not just an add-on, as you will learn when you start customizing menus. In fact, all of the main programs are primarily built using ELF.

Applix has made a serious commitment to Java. Most of the client stuff is available in Java thin-client versions under the name Applix AnyWare. For the purposes of this review, I have concentrated on Applixware, which is the UNIX desktop office suite. However, even Applixware uses Java. Some of the import filters, for example, are written in Java. This is quite sensible because it simplifies porting, but doesn't have a serious performance penalty because the importing is typically done less frequently than basic operations. The speed of the import operations was quite acceptable.

Invocation

While you can start Applixware with no arguments, you can also invoke it with a file type and/or file name. There are options like -wp, which tells the program you want to do word processing. These can be combined with a file name that will be opened or imported into a new document, depending on whether or not it is an Applix format. Using any

Figure 3. Building a GUI

of the UNIX desktops—CDE, KDE and so on—it would be possible to bind these types for click-to-start. Applix arguably should do this for you. It doesn't. For now, I have written a shell script (see Listing 2, Page 70), similar to the Windows 95 start command, because I use a variety of window managers and desktops. That this invocation works so well is a testament to Applix's understanding of how UNIX programs ought to behave.

Applix Improvable

I now use Applix a lot, even though I'm not fully satisfied with it. The main shortcomings I have encountered are in the import filters and the GUI. First, there are no "tool tips" on any of the toolbars. Tool tips were first shown as Balloon Help on Apple Computer Inc. Macintosh systems, but in its usual style of steal-and-improve, Microsoft brought them online as ToolTips. ToolTips don't pop up obtrusively the nanosecond you visit a toolbar icon—this was one problem of Apple's original—instead, they wait until you hesitate with the mouse

Listing 2. Start Shell Script

```
#!/bin/ksh -f

# simple emulation of Windows' command of the same name:
# deduce type of file, start its application (backgrounded).

for f
do
        case ${f} in
            *.ai)   illustrator ${f} & ;;
            *.aw|*.doc|*.wpd)       applix -wp ${f} & ;;
            *.as)   applix -sc ${f} & ;;
            *.class) java ${f}& ;;
            *.fm)   appplix -wp ${f} & ;;
            ftp://*) ftp ${f} & ;;
            *.gif)  xv ${f} & ;;
            http:*|*.htm|*.html) netscape -remote "openURL(file:'pwd'/${f} )"& ;;
            *.pdf)  acroread ${f} & ;;
            *.ps)   gs ${f} & ;;
            *.txt)  dtpad ${f} & ;;
            *.xls)  applix ${f} & ;;
            *)              echo "Duh, I dunno how to start ${f}!!" >&2;;
        esac
        sleep 5         # to avoid killing the machine
done
```

over a tool for some reasonable fraction of a second. Even my 15-line Swing-based Java program can have tool tips. Why is it so hard here? Is it a Motif limitation that really can't be worked around? Or is it merely a lack of time? Whatever the case, the lack of tool tips

does make it harder to learn a new GUI tool. No matter how intuitive a designer thinks his "universal icons" to be, they aren't. A well-chosen phrase is worth a thousand pictures that don't represent anything.

Speaking of the GUI, there are also too many dialogs that don't learn the directory you're in from each other. You have to navigate around to some directory Open and then Import something, for example, and the Import has no recollection of the directory you opened from; yet most of the time, I want to do things from the same directory.

A minor glitch of the GUI code is it reverts to black and white if something about your display isn't as expected. It did this a few times when I thought nothing had changed. Perhaps it is merely susceptible to running out of colors; because my IPX has only a CGSIX display, its 256 colors can be gone in a flash if I'm running a Web browser.

One more pet peeve: the default menu shortcuts. Under the File menu, ^S for save and ^P for print are familiar, and imply Mac/Windows compatibility, but then F2, F3 and so on are used for edit/cut/paste, undo/redo and so forth. Someone used to conventional shortcuts would have expected ^C, ^X, ^V and ^Z for these, given the use of ^S and ^P.

It turns out to be fairly easy to customize, but it's hidden beneath the Customize menu bar. It was almost no work to add this same functionality under the more traditional Edit-Preferences menu bar, using the Customize menu bar feature. Next, I wanted to change the menu item accelerators to ^C, ^X and so on. Unfortunately, ^X was already defined as some delete key function. And it became tedious to fix several such duplicates because on each one, dismissing the Error dialog also dismissed the menu bar dialog, which is inconsistent with any other GUI program. When it did this, it discarded all my changes.

I still wanted to use ^X and ^W, but these conflicted with key assignments. I couldn't find a "Customize Keys" feature, and there was nothing in the on-line help. Eventually, I was able to edit them using the ax_wp4 file in vi.

Overall, Applixware is a good office suite for UNIX, and has all the tools you need. It's strong on functionality, with a few GUI sections needing refinement and a few wrinkles in the import filters. ➥

Ian Darwin is a consultant, developer and trainer specializing in UNIX operating systems and the Java language. Email: ian@darwinsys.com

Applixware 4.41

Company
Applix Inc.
112 Turnpike Road
Westboro, MA 01581

Phone
(508) 870-0300

WWW
http://www.applix.com

Best Feature
Full-function office suite native to UNIX. ASCII save files.

Worst Feature
Price gap between commercial UNIX (Solaris) and Linux versions. Limited word processing "named" formats and a few problems importing some file formats.

Price
Applixware (Solaris version) $495
Applixware for Linux $99

Circle 160

NEW PRODUCTS

Fibre Channel-to-SCSI Bridge Unveiled

Atto Technology is now offering Atto FibreBridge Enterprise, a Fibre Channel-to-SCSI bridge. With FibreBridge, users can attach any SCSI-based device to Fibre Channel, the company says. The product comes with dual Gigabit Interface Converters (GBICs) and features Ultra 2 Low-Voltage Differential (LVD) and High-Voltage Differential (HVD) SCSI bus compatibility. In addition, FibreBridge supports the Storage Network Industry Association (SNIA) Extended Copy command. FibreBridge uses SNIA Extended Copy, along with Legato Systems Inc.'s Celestra—an imbedded system for running "serverless" backup, file replication and mirroring—to

run backup without placing load on the server. Data is backed up directly from a disk array to a tape library without putting extra load on the server, Atto says. This is said to remove backup traffic from the local area network (LAN) and improve performance access to data and applications. Contact vendor for pricing.

Atto Technology Inc.
40 Hazelwood Drive, Bldg. 106
Amherst, NY 14228
http://www.attotech.com
Circle 101

Continuous Online UPS Protection

Tripp Lite's new 6KVA and 10KVA Unison On-Line UPS systems continuously convert incoming AC power to filtered DC power, then resynthesize it to low-harmonic, frequency-controlled AC power, the company says. This constant online operation generates a pure sine wave output with zero transfer time, while completely isolating connected

equipment against power problems such as blackouts, surges and line noise, Tripp Lite says. A unique feature of both systems is a Line-Interactive Mode option, which allows users to switch the UPS to an energy-saving mode when online protection isn't needed. This feature is aimed at increasing the UPS system's overall power efficiency, resulting in reduced operating costs without affecting system output reliability under brownout or blackout conditions.

Each product comes with Tripp Lite's PowerAlert Software, allowing users to manage and control all UPS systems

Entry-Level SPARC Gets 440-MHz Power

Tatung has unveiled its fastest entry-level SPARC system to date, the COMPstation U10-440. Tatung says the single-processor U10-440, which is available in tower and rack-mount configurations, takes advantage of the recently released 440-MHz UltraSPARC-IIi processor from Sun Microsystems Inc.

Equipped with 2 MB of external cache, the 64-bit 440-MHz processor can support up to four 32-bit PCI devices at 33-MHz and features a 64-bit UPA slot for vertical add-ons such as Creator3D and Elite3D graphics cards. While incorporating the PCI I/O bus to provide users with access to a host of peripherals, the U10-440 also maintains full binary compatibility with existing software on SBus-based systems, Tatung says.

The COMPstation U10-440 comes equipped with five drive bays that can be fitted with two 3.5-inch hard drives, one 5.25-inch CD-ROM, one 4- or 8mm tape drive and either one 3.5-inch floppy or another peripheral. Standard features include 128 MB of RAM, 9-GB hard drive, one PCI graphics card, one 64-bit UPA slot, two serial ports, one parallel port, dual-channel Ultra Wide SCSI on the motherboard, 10/100-BaseT Ethernet interface and Solaris 7 preinstalled.

Optional features include 1 GB of memory, 24-bit PCI graphics, Creator3D, Elite3D m3 or m6 graphics cards, Asynchronous Transfer Mode (ATM) connectivity and an

Ultra Wide SCSI adapter. The company is also offering optional monitor sizes for the U10-440, including 17-, 19-, 21- and 24-inch monitors. Pricing starts at $4,690 for a standard tower configuration and $4,780 for a standard rack-mount configuration.

Tatung Science & Technology Inc.
1840 McCarthy Blvd.
Milpitas, CA 95035
http://www.tsti.com
Circle 100

throughout an enterprise. PowerAlert provides a single GUI of UPS systems networkwide and lets users act on alerts remotely without having to purchase an expensive network management system, Tripp Lite says. The Unison 6KVA and 10KVA UPS systems cost $5,499 and $9,499, respectively.

Tripp Lite
1111 W. 35th St.
Chicago, IL 60609
http://www.tripplite.com
Circle 102

Appliance Offers Flexible Data Mirroring

Ark Research, a provider of mission-critical business solutions, has announced Ark/2000. Ark Research calls Ark/2000 the first flexible, cost-effective and highly scalable server appliance for mirroring data to storage devices located anywhere throughout an enterprise and/or wide area network (WAN).

Designed for today's heterogeneous environments, Ark/2000 reportedly offers nondisruptive remote/local data mirroring, continuous backup, data warehousing and rapid disaster recovery, while also providing improved performance for data access and file transfer. Ark Research says it is the only storage solution to provide real-time, platform- and operating system-independent data duplication and access over a WAN.

Once implemented, updates to files, failover operations and backup file access are automated. Ark/2000 can be easily attached to any file server-based network and installed between the host computer and any mass storage device using SCSI connections. The product's design is said to permit the interconnection of multiple units via Fast Ethernet, T1 or T3 links. Two or more Ark/2000 units on a network provide users with a robust, centralized disaster recovery system, the company says.

Ark/2000 is priced at $14,995 per unit and comes with an Intel Corp. Pentium II processor, 256 MB of cache, four SCSI 3 interfaces (Fast, Wide, Ultra and Ultra 2), two 10/100-Mb/s Ethernet ports, two serial interfaces for remote control and monitoring, and an internal hard disk.

Ark Research Corp.
1190 Saratoga Ave., Ste. 110
San Jose, CA 95129
http://www.arkres.com
Circle 103

Content Management for Solaris

Eprise has added support for Sun Microsystems Inc.'s Solaris platform to its content management application, Participant Server 2.0. The company has also added Java support, which allows Participant Server to integrate with Java-based application servers.

Participant Server is a system for business professionals to contribute and manage information via a Web site, the company says. With Participant Server, authorized users can independently edit content using any Java-enabled Web browser. The added Java support includes a JavaBeans interface, which connects to the server's eXtensible Markup Language (XML) integration framework. Also, the Java support extends the content management product to Java-Server Pages (JSP), Java servlets and other Java applications, Eprise says.

Pricing for Participant Server 2.0 starts at $50,000; the JavaBeans interface costs $5,000. It runs on Solaris or Windows NT.

Eprise Corp.
1671 Worcester Road
Framingham, MA 01701
http://www.eprise.com
Circle 104

Java Source Code Tool Upgrade

TakeFive Software has released SNiFF+J 3.1, the latest version of its Java technology source code tool for multiplatform development. With SNiFF+J, developers working in multi-language client/server environments can navigate, analyze and manage a variety of source code. According to TakeFive,

SNiFF+J can handle millions of C++ and Java lines while also integrating with other development tools, including source code management tools, CASE tools, editors and debuggers. SNiFF+J 3.1 can generate Java code specifically for Solaris, HP-UX, Linux, IRIX, AIX and Windows 95/98/NT platforms. A single license costs $1,750.

TakeFive Software Inc.
20813 Stevens Creek Blvd., Ste. 200
Cupertino, CA 95014
http://www.takefive.com
Circle 105

3D Graphics Workstations Unveiled

Hewlett-Packard's Visualize B1000, C3000 and J5000 PA-8500 RISC-based systems are said to offer new workstation options for high-end graphics, engineering and technical customers.

HP Visualize B1000 for 2D and 3D design is aimed at CAD designers and 3D modelers who design subassemblies and components. It comes with a 300-MHz PA-8500 processor, 1.5-MB on-chip cache, 128-MB to 4-GB RAM and six PCI slots. The Visualize C3000 is touted as an ideal platform for mechanical engineers performing simulations, virtual prototyping, complex modeling or high-end visualization. It comes with a 400-MHz PA-8500 processor, 1.5-MB on-chip cache, 512-MB to 4-GB RAM and eight PCI slots.

Finally, the dual-processor J5000 is well-suited for compute-intensive IC verification and simulation compute farms in electronics and advanced 3D design, HP says. It handles complicated computational analysis and computer- and memory-intensive processing work-

loads. The J5000 comes with two 440-MHz PA-8500 processors, 1.5-MB on-chip cache, 512-MB to 5-GB RAM and eight PCI slots.

Pricing starts at $9,888 for the B1000, $13,763 for the C3000 and $22,636 for the J5000.

Hewlett-Packard Co.
3000 Hanover St.
Palo Alto, CA 94303
http://www.hp.com
Circle 106

App Server Supports Java 2

SilverStream Software's SilverStream 2.5 application server now includes support for the Java 2-compliant Java Virtual Machine from Sun Microsystems Inc., as well as enhanced support for the latest databases, such as the IBM Corp. DB2 5.2 and Microsoft Corp. SQL Server 7.0, and a native Java database connectivity (JDBC) driver for Informix Corp. databases. SilverStream 2.5 also features improved wizards for linking applications to data source objects and Korean language translation.

The SilverStream 2.5 application server runs on Solaris, HP-UX and Windows NT and sells for $10,000 for a single-CPU system. A Group Developer Pack, with five programmer licenses and a five-user test application server, costs $2,500.

SilverStream Software
1 Burlington Woods
Burlington, MA 01803
http://www.silverstream.com
Circle 107

New Library to Improve STEP Interoperability

Step Tools, a developer of STEP (Standard for the Exchange of Product Model Data) integration software tools based on the International Organization for Standardization (ISO) standard, has released ST-Parasolid Library, a C++ library that is used to build customized STEP translators for products based on the Parasolid solid modeling kernel.

The library is said to deliver state-of-the-art programming and desktop tools for translating Parasolid solid models used in CAD/CAM/CAE applications. Developers can reportedly use the library

to create a custom bidirectional translator that will map STEP geometry and topology directly onto Parasolid models, without the use of an intermediate file format. According to the company, MCAD users will now be able to export Parasolid models via STEP to non-Parasolid systems. The library can also be used to implement solid modeling operations on STEP data.

ST-Parasolid Library comprises a set of C functions that can be used to create STEP files that are compliant with AP-203 and AP-214 application protocols. In addition, it comes complete with source code and a sample translator, and

is available for Solaris, AIX, IRIX, HP-UX and Windows NT platforms. Training, consulting and technical support programs are also available. A single license for ST-Parasolid Library costs $15,875; a full source code license is also available for $81,250. (Note: This product requires ST-Developer and Parasolid licenses from Unigraphics Inc.)

Step Tools Inc.
1223 Peoples Ave.
Troy, NY 12180
http://www.steptools.com
Circle 108

New Optimization Module for DesignMaxx

Opmaxx has introduced a mixed-signal optimization module for its DesignMaxx tool, one of a family of mixed-signal design and test automation tools. With the new optimization module, analog and digital circuit designers can reportedly automate the process of establishing design parameters for circuits.

The DesignMaxx optimization module makes optimization practical for most analog and high-speed digital circuits, the company says. On the analog side, this includes amplifiers, filters, converters and timers; on the digital side, circuits include high-speed data paths and I/O buffers.

DesignMaxx can determine the gradient of the output specification with respect to every circuit parameter of interest, the company says. The optimization module can then calculate the specification of a mixed-signal design that can include timing, power consumption, total harmonic distortion, pole-zero, integral nonlinear, differential nonlinear and unity-gain bandwidth. Design restraints can then be added to ensure parameter optimization is complete, Opmaxx says.

DesignMaxx is available for Solaris and HP-UX. The optimization module costs $40,000 and DesignMaxx costs $48,000.

Opmaxx Inc.
8700 S.W. Creekside Place
Beaverton, OR 97008
http://www.opmaxx.com
Circle 109

Fault-Tolerant App Server Software

The ObjectSwitch 3 product family from ObjectSwitch provides software fault tolerance, hot-swapping of online applications, lower development costs for new adapters and interfaces and support for open Unified Modeling Language (UML)-based object modeling environments, the company says.

ObjectSwitch 3 is designed to offer improved scalability and high-speed transactions to telecommunications and networking companies. The product

Upgrades, Enhancements, Additions...

■ Enlighten Software Solutions is extending the platform support of its flagship product, EnlightenDSM, to include Linux. EnlightenDSM is a standards-based, multifunction system for managing mixed UNIX and Windows NT environments via a central console. The product already supports most flavors of UNIX, including SunOS, Solaris, AIX, HP-UX and IRIX, and Windows NT. This move, the company says, will extend the acceptance and usability of Linux by providing a cost-effective, easy-to-use suite of mulitplatform systems management tools to facilitate Linux management. Enlighten Software attributes this decision to heightened customer interest in the open-source operating system. The Linux version of EnlightenDSM is slated for release by the end of third-quarter 1999. Contact company for pricing. **Enlighten Software Solutions Inc.**, 999 Baker Way, 5th Floor, San Mateo, CA 94404, *http://www.enlightendsm.com*. **Circle 110**

■ Compatible Systems' IntraPort Release 3 client software for virtual private networks (VPNs) now uses the industry standard Internet Key Exchange (IKE) protocol to establish and manage a VPN client session. Prior versions used a proprietary protocol. IKE greatly simplifies the authentication process, the company says. IntraPort client software works with the IntraPort family of VPN access servers and runs on Solaris, Linux, Mac OS and Windows 95/98/NT. **Compatible Systems Corp.**, P.O. Box 17220, Boulder, CO 80308, *http://www.compatible.com*. **Circle 111**

■ XIOtech has announced extended operating system support for Magnitude, its Storage Area Network (SAN) in a box. Magnitude uses the company's Real-time Data Intelligence (REDI) storage processing architecture and offers up to 3.2 TB of centralized storage via a new Fibre Channel Arbitrated Loop (FC-AL) server connection. This new connection allows Magnitude to connect to up to 192 servers, each running one of a number of supported operating systems, including Solaris, AIX, IRIX, HP-UX, Linux, NetWare and Windows NT, the company says. Previously, the storage system's point-to-point connection allowed for only eight servers running either NetWare or Windows NT to connect. With FC-AL, servers can connect to Magnitude in either a point-to-point or a loop topology, XIOtech says. The FC-AL connection can be installed in all models of Magnitude and requires users to install a XIOtech FC-AL host adapter board in one of Magnitude's eight PCI ports and a third-party FC-AL adapter and operating system driver in the attached server. Pricing for Magnitude costs between $.09 and $.30 per Mb, depending on configuration. **XIOtech Corp.**, 6509 Flying Cloud Drive, Ste. 200, Eden Prairie, MN 55344, *http://www.xiotech.com*. **Circle 112**

■ Chili!Soft ASP 3.0 now supports IBM Corp.'s AIX operating system and the open-source Apache Web server. Chili!Soft ASP expands the platform independence of Microsoft Corp.'s Active Server Pages (ASP) technology for creating dynamic Web sites, the company says. In addition, the product now ships with Chili!Beans, a Component Object Model (COM)-to-Java bridge, which allows ASP developers to access Java classes using ASP scripts. ASP developers can connect to Sun Microsystems Inc. Enterprise JavaBeans (EJB) running in an EJB application server, or use Java 2's Common Object Request Broker Architecture (CORBA) support to connect to legacy CORBA objects. Additional features include a choice of security modes, new configuration options and an improved engine for faster, more scalable performance, the company says. **Chili!Soft Inc.**, 13920 S.E. Eastgate Way, Ste. 120, Bellevue, WA 98005, *http://www.chilisoft.com*. **Circle 113**

■ Tenon Intersystems has upgraded its XTen X Window server for Apple Computer Inc.'s Power Macintosh platform to include a fast native file system, improved X Display Manager (XDM) session control, secure shell remote client launch, multiple password-based user desktops, hot-key `xterm` launch, integrated Common Desktop Environment (CDE) fonts and support for shared memory extensions. XTen supports the X Display Management Control Protocol (XDMCP) and can be configured as an XDM manager. It can also be configured as a font server or used as a remote font server. For added compatibility, XTen now uses Open Transport as a default, and also includes Tenon's multilink TCP protocol stack. XTen costs $175 per single copy. **Tenon Intersystems**, 1123 Chapala St., Santa Barbara, CA 93101, *http://www.tenon.com*. **Circle 114**

family comprises ObjectSwitch Telecom Integration Server, ObjectSwitch Transaction Application Server, ObjectSwitch Design Center IDE (integrated development environment) and Object Model Auditor. The ObjectSwitch 3's memory-resident and network-resident architecture enables application autorecovery with no single point of failure, the company says.

The software runs on Solaris, HP-UX and Windows NT. Pricing for the ObjectSwitch 3 product family is as follows: Telecom Integration Server costs $75,000; Transaction Application Server costs $75,000; Design Center IDE costs $2,500; and Object Model Auditor costs $1,000.

> **ObjectSwitch Corp.**
> 900 Larkspur Landing Circle
> Ste. 270
> Larkspur, CA 94939
> *http://www.objectswitch.com*
> **Circle 115**

Interoperability Product Line Enhanced

Mortice Kern Systems (MKS) has enhanced three offerings from its line of UNIX and Windows interoperability products: MKS Toolkit 6.2, MKS Toolkit Select 6.2 and MKS NuTcracker Professional 4.2.

MKS Toolkit 6.2 is said to provide a UNIX command-line environment on Windows that includes a suite of more than 350 UNIX and Windows NT utilities. Developers and systems administrators can customize, control and automate software administrative processes, the company says. The 6.2 release now includes MKS Cshell and KornShell. Users can choose which command line and scripting environment to use on Windows 95/98/NT. MKS Toolkit Select 6.2 addresses the interoperability issue on a platform level, the company says. It provides access to MKS Toolkit utilities and interactive access to remote UNIX applications from Windows clients.

MKS NuTcracker is designed to manage application portability by providing a system for building and deploying character- or GUI-based UNIX applications on Windows. NuTcracker can be used to develop single- or multi-threaded C, C++ or FORTRAN applica-

tions, shared libraries and daemons on Windows 95/98/NT. MKS says these applications can maintain a common source code baseline across Windows, UNIX and Linux.

MKS Toolkit 6.2 and MKS Toolkit Select 6.2 are available for a suggested retail price of $399 and $549 per single user, respectively. MKS NuTcracker Professional 4.2 costs $5,000.

> **Mortice Kern Systems Inc.**
> 185 Columbia St. W.
> Waterloo, Ontario
> Canada N2L 5Z5
> *http://www.mks.com*
> **Circle 116**

Copper-to-Fiber Media Converters Out

New copper-to-fiber media converters for 10/100-Mb/s full-duplex and half-duplex Ethernet local area networks (LANs) and wide area networks (WANs) are now available from IMC Networks. The McBasic 10/100-Mb/s Ethernet media converter and the McV 10/100

Module are said to allow network managers to connect 10/100-Mb/s copper-based hubs, switches and routers with fiber-based cabling infrastructures.

The McBasic is a stand-alone product designed for small companies needing to connect remote offices and comes with a 100/240-VAC internal power supply.

The McV 10/100 Module installs into an IMC Network McV/21x Series chassis, which can host up to 21 front-loading, hot-swappable modules alongside redundant, modular power supplies. The McV 10/100 Series is designed for larger installations that have extensive fiber requirements, IMC says.

The McBasic 10/100 costs between $450 and $500, depending on configuration, and the McV/21x Chassis costs $650. Pricing for the McV 10/100 Module has not been determined.

> **IMC Networks**
> 19772 Pauling
> Foothill Ranch, CA 92610
> *http://www.imcnetworks.com*
> **Circle 117**

QUICK RESPONSE TO YOUR INFORMATION REQUESTS

For more information on the products/services advertised in this section, please circle the appropriate reader service number(s) on the reader inquiry card located elsewhere in this issue. OR

You can E-mail your requests for more information using the reader inquiry card numbers.

Here's how:

E-mail address: **BERKCOMP@AOL.COM**

Include

1. Magazine name and issue date
2. Your name and address
3. Reader inquiry numbers, separated by commas, for which you need more information.

Server/Workstation Marketplace

86

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 St., Ste. 9
North Grafton, MA 01536
Phone: (508) 839-4016
Fax (508) 839-4226
Email: caflanag@aol.com

New York/Mid-Atlantic/Southeast

JOANNA PARKER
18 Stephanie Drive, Ste. 3
Stirling, NJ 07980
Phone: (908) 542-0789
Fax: (908) 542-0782
Email: joanna@cpg.com

Mid-West/Southwest/Central Canada

LINDA LIEBICH
9600 Great Hills Trail, Ste. 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 Ave.
San Jose, CA 95130
Phone: (408) 374-9925
Fax: (408) 374-9926
Email: vmiyaoka@cpg.com

Server/Workstation Marketplace Classifieds/Recruitment Ads

CAROL FLANAGAN Mgr, Telemarketing Sales
Phone: (508) 839-4016
MARY ERICSON
Phone: (508) 839-0720
80 Worcester St., Ste. 9
North Grafton, MA 01536
Fax: (508) 839-4226
Email: caflanag@aol.com

ALL INTERNET. ALL BUSINESS. ALL WEEK.

From advertising to e-commerce to enterprise connectivity, when it comes to putting the power of the Internet to work for business, only one event puts all the possibilities at your fingertips — Fall Internet World.

Our five-day conference is the best in the industry — covering all the Internet business applications, technologies and trends your company needs to increase performance.

Then there's our exhibit hall — the world's largest display of Internet solutions from more than 400 leading industry suppliers, including Microsoft, IBM, HP, AT&T, Oracle, RealNetworks — all the big names.

If the Internet is important to your business, come to the event that's all Internet, all business, all week.

OUR 1999 KEYNOTES:

Craig R. Barrett
President/CEO
Intel Corporation

Larry Ellison
Chairman/CEO
Oracle Corporation

Ivan Seidenberg
Chairman/CEO
Bell Atlantic Corporation

Sanjay Kumar
President/COO
Computer Associates
International, Inc.

Joseph P. Nacchio
Chairman/CEO
Qwest Communications
International Inc.

**Register online now and SAVE! $50 off exhibit hall admission or $100 off the conference
http://www.internet.com/registerfall – Discounts expire Sept 17!**

For complete event details:
Visit:
http://events.internet.com/fall99
Call: 1-800-500-1959
E-mail: fiwprogram@iw.com

To exhibit:
Call: (203) 341-2969
E-mail: smoriarty@iw.com

THE SEVENTH ANNUAL
FALL Internet World 99
CONFERENCE: OCT. 4–8
EXHIBITS: OCT. 6–8
JACOB K. JAVITS
CONVENTION CENTER
NEW YORK, NY

Produced by:
Penton

Media Sponsors:
internet.com

Corporate Sponsor:
intel.

Official Sponsor:
internet WORLD

boardwatch

Show Network Provided by:
**IBM** From the E-mail Center to the show floor... Are you networked for e-business?

Code: N57

The
ANATOMY
## Of An Enterprise-Class Library

I. The Brain – *Much like "some" human brains, the P3000 has a massive capacity to store and move information. This intelligent library has a native capacity of 11.4 terabytes and blazing performance of 288 gigabytes per hour.*

II. The Eyes – *With local and remote browser GUIs, you'll see the industry's most powerful DLTtape library is incredibly easy to use.*

III. The Skeleton – *The human body has two arms and two legs. The P3000 delivers the same high availability (HA) design with redundant AC cords, power supplies and fans. Plus, the power supplies, fans and DLTtape drives can be hot-swapped.*

VII. The Spine – *The backbone of the P3000's design is a PCI expansion bus supporting SCSI interface, Fibre Channel, tape array and server PCI cards – "future proofing" your library with a modular upgrade path.*

Cervical
Thoracic
Lumbar
Sacral
Coccyx

Divisions of the Spine

IV. The Nervous System – *The complex nervous system of the P3000 is designed to support multiple concurrent network, SCSI and fibre channel connections, so each library can be shared by NAS, SAN and direct-connect environments.*

VI. The Heart – *The heart of the P3000 is the IntelliGrip precision cartridge handling system which will pick-and-place cartridges for years without skipping a beat.*

V. *Like a well-tuned body, The P3000's reliability, redundancy, ease of use and modular upgrades all add up to low total cost of ownership (TCO).*

ATL
PRODUCTS
A **Quantum** Company

DLT TAPE

DLTTAPE LIBRARIES DESKTOP TO DATACENTER

Circle No. 23