

UNIXTM WORLD

THE MAGAZINE FOR MULTIUSER, MULTITASKING SYSTEMS

DECEMBER 1985

\$3 IN USA

£2.20 IN UK

A TECH VALLEY PUBLICATION




REVIEW:
AT&T
System V
Training

Sun's NFS
& AT&T's RFS



SNA, SNADS, XNS,
TCP/IP (And More)
Defined

Inside Edge:
Alliant's "Crayette"

Practical Usenet,
UUCP Applications



LOCAL-AREA
NETWORKING



UNIX/WORLD

THE MAGAZINE FOR MULTIUSER, MULTITASKING SYSTEMS

VOLUME II, NUMBER 11

C O N T E N T S

A TECH VALLEY PUBLICATION

THEME

WAITING FOR BIG BLUE GODOT *by Michele Nielson* An honest evaluation of the sorry state of today's local-area networking marketplace.

30


SUN's NFS AND AT&T's STREAMS *by Mark Hatch* Beyond local-area networks lies an as yet unconquered territory—heterogeneous distributed file systems. Here's a look at two possible solutions to the distributed file system quandry no local-area networking vendor has told you about yet.

38

DECIPHERING DATA COMMUNICATIONS TERMINOLOGY

by Judy Estrin Ever wonder what techno-babble such as SNA, SNADS, XNS, TCP/IP, and company really mean? A noted authority on local-area networking technologies clears the air.

54



30

REVIEW

AT&T's UNIX SYSTEM V TRAINING *by Paul Fertal* UNIX/WORLD's own intrepid system administrator ventured forth to the Unix-Versity to learn the deep, dark secrets of the Unix-Verse. Here is his report.

65

JOURNAL

DOING RESEARCH ON USENET *by Gary Perlman* There are practical aspects to Usenet, the semi-underground network that links Unix system users worldwide. Our author explains.

75



38

TRENDS

EDITOR'S CONSOLE 7
FOR THE RECORD 9
mail 15
INSIDE EDGE 17
NEW PRODUCTS 96
sync 110

TRAINING

WIZARD'S GRABBAG 82
STARTER KIT 86
/usr/LIBRARY 89
CALENDAR 106

/etc

ADVERTISERS' INDEX 90
MARKET PLACE 106
CAREER OPPORTUNITY

NEXT MONTH

NEXT MONTH:
■ **TECHNOLOGY/
MARKET
FORECASTS**
■ **THE BEST OF
UNIX/WORLD**

UNIX/WORLD (ISSN 0739-5922) is published monthly by Tech Valley Publishing, with offices at 444 Castro St., Mountain View, CA 94041, telephone 415/940-1500. RETURN POSTAGE GUARANTEED. Postmaster, send Form 3579 to UNIX/WORLD, P.O. Box 1929, Maron, OH 43305.

Subscriptions cost \$18 for one year within the U.S.A. and its possessions. For Canadian subscriptions, add \$6 for each year. For foreign subscriptions, add \$20 each year for surface mail. Remit in U.S. funds drawn on a U.S. bank only. *Subscription questions and problems:* Contact UNIX/WORLD Subscription Department, P.O. Box 1929, Maron, OH 43305. Printed in United States of America. Requests for reprints or bulk orders should be addressed to the Merchandising Manager.

Address all editorial correspondence to the editor at UNIX/WORLD, 444 Castro St., Mountain View, CA 94041. Unsolicited manuscripts can be returned only if accompanied by a self-addressed, stamped envelope. UNIX/WORLD is not responsible for lost or damaged manuscripts, photos, or artwork. Entire contents copyright © 1985 by Tech Valley Publishing, unless otherwise noted. All rights reserved.

UNIX™ is a trademark of AT&T Bell Laboratories. WE and UNIX System III are trademarks of Western Electric. UNIX/WORLD is not affiliated with AT&T Bell Laboratories or Western Electric. Apollo, DOMAIN, and DOMAIN/IX are trademarks of Apollo Computer Inc. DEC, VAX, PDP, VMS, Ultrix-32, and DEC Rainbow are trademarks of Digital Equipment Corp. IBM, PC, PC/AT, PC/IX, and PC-DOS are trademarks of International Business Machines Corp. XENIX, MS-DOS, and Multiplan are trademarks of Microsoft Corp. CP/M, CP/M-86, GEM, and CP/M-68K are trademarks of Digital Research Inc. Ethernet is a trademark of Xerox Corp. M68000 is a trademark of Motorola Inc. HP and HP-UX are trademarks of Hewlett-Packard Co. SUN Microsystems and NFS are trademarks of SUN Microsystems Inc. UniFlex is a trademark of Technical Systems Consultants. UNOS is a trademark of Charles River Data Systems. VENIX is a trademark of Venturcom Inc. Idris is a trademark of Whitesmiths Ltd. Coherent is a trademark of Mark Williams Co.

Statements of fact and opinion are made on the responsibility of the authors alone and do not imply an opinion on the part of Tech Valley Publishing or the editorial staff of UNIX/WORLD.

USENET: DOING RESEARCH ON THE NETWORK

The author describes his experiences while doing research using the Unix system network Usenet.

BY GARY PERLMAN

An interesting feature of the Unix system world is the Unix system user network, Usenet. Usenet allows users to post articles to distributed bulletin boards, collectively called *netnews*, that are read by thousands of technical people every day. There are bulletin boards for every subject for which there is sufficient interest, such as political groups, cooking groups, groups for telling jokes, groups for letting off steam, and even some groups related to research. In particular, there are special-interest groups that cover Unix system topics.

The distribution mechanism for Usenet is a prime example of technical people's distaste for protocol. To receive netnews, a new machine has to get the cooperation of a neighboring Unix system machine that will feed news articles to the new site and distribute articles from the site. Subscription to news groups is on a per user basis. Distribution takes from a few minutes for direct connections to several days for connections through dozens of machines. Articles travel all over the world—to Europe, Australia, and some to

Southeast Asia. The result is a network of about 2000 machines (Summers-Horton and Horton 1985), each with a unique Usenet connection to any other machine. Note that this is independent of the UUCP (Unix to Unix CoPy) network in which each machine is connected to several others; this network is used for electronic mail and file transfers (see sidebar, *UUCP* and *Usenet* explained).

The content of most articles on Usenet is completely unmoderated; although a few moderated groups post summaries of discussions. In part, the lack of moderation was blamed for the unauthorized distribution of a copyrighted paper criticizing the Unix system user interface (Norman 1981). This event clarified the network's position as a communication tool comparable to television for use and abuse. The only type of editorial action that occurs on the network is that some sites do not distribute certain groups to certain sites—for example, many sites do not carry jokes, and some companies keep internal news groups that are not distributed outside the company.

A particularly useful aspect of the Unix system network is that it can facilitate research in ways never before possible. Think of it—2000 machines all running the same operating system, each with dozens of users who have access to the same software. For someone interested in research on the Unix system, easy access to most of the Unix system user community is a dream come true. You can distribute prototype versions of programs, drafts of papers, and requests for information to special-interest groups, for example.

In the following sections, I will describe some of my experiences with different types of research using netnews on Usenet, or “the net” for short.

GATHERING SIMPLE INFORMATION

A good example of how you can use the net to gather information from many sources with little effort is to use the netnews group `net.wanted`, a bulletin board to which people seeking information post requests. I decided it would be fun to be able to spell out words using radio jargon—such as NOVEMBER, ECHO, and TANGO for spelling N-E-T. I posted a request asking people to send responses directly to me and said I would post a summary on the net. This is customary practice because if everyone posted their replies to the net, it would only add to the already heavy traffic.

Still, some people did post their replies directly to the net—I suspect they wanted recognition for their efforts (see the later discussion about recognition). I received about ten replies and was able to post a summary that well exceeded my original expectations. To my

discredit, I did not acknowledge contributors.

GATHERING STRATEGIC INFORMATION

Over my years of Unix system use, I have developed a public-domain statistical package (Perlman 1984). This prompted me to post a request for a list of programs or program features people wanted. I knew the answers would be biased because of the characteristics of people on the net and of the subset who reply to requests.

I received about 50 replies, which I abbreviated and posted. I did not get much useful information—although it was interesting to see a "wish list" from Unix system programmers. On the other hand, I did not expend much effort to get it, and it was of interest to marketing people.

SPYING ON THE NET

After getting information from the net, I thought it would be fair to provide some of my own. I decided to post the abstract of some papers of

mine to the net so that people could request the full papers. For this, I made use of subject-specific news groups such as `net.cog-eng` (cognitive engineering or human factors) and `net.text` (document preparation). I have found this to be an effective way to get people to request technical report—I got more requests from netnews readers than from the readers of the journals the papers were published in!

I kept a list of the people requesting papers so that I could send them new papers on the same topic, and while doing this noticed some

The following is a reformatted version of the list of active Usenet newsgroups (as of April 2, 1985) posted by Gene Spafford. `spaff@gatech` (netnews message id 12758@gatech.CSNET).

Usenet postings are made to newsgroups and are restricted by a distribution specification. There are three newsgroups for the whole net:

`net` newsgroups circulated to the entire network
`fa` primarily digests From the Arpa net
`mod` moderated groups

Distributions can also be restricted within geographical or corporate boundaries.

`net` worldwide distribution (the default)
`atl` limited to AT&T
`can` limited to Canada
`eunet` limited to European sites
`na` limited to North America
`usa` limited to the United States

Most U.S. states have distribution categories named after the two-letter abbreviation for that state or category (that is, `ga` for Georgia, `nj` for New Jersey, and so on).

What follows is a long list of netnews groups with descriptions (for obvious cases, no description is provided). Specialized subgroups are indented recursively; some are listed all in one or two lines. To form a newsgroup name you begin with a distribution (such as `net`) and add the names of special interest groups and sub-groups (that is, `math` and `stat`) with periods (called "dots") in between (that is, `net.math.stat`).

Groups	Descriptions
<code>net</code>	all netnews groups
<code>abortion</code>	
<code>ai</code>	artificial intelligence
<code>analog</code>	analog design developments, ideas, and components
<code>announce</code>	moderated, general announcements of interest to all
<code>newusers</code>	moderated explanatory postings for new users
<code>arch</code>	computer architecture
<code>astro</code>	astronomy
<code>expert</code>	experts in astronomy
<code>audio</code>	high fidelity audio
<code>auto</code>	automotive products and laws
<code>aviation</code>	rules, means, and methods
<code>bicycle</code>	products and laws
<code>bio</code>	biology
<code>books</code>	all types of books
<code>bug5</code>	general bug reports and fixes for Unix 2bsd and 4bsd (Berkeley Unix systems), u5g (Unix Systems III, V), uucp, v7
<code>chess</code>	chess and computer chess
<code>cog-eng</code>	cognitive engineering and user interfaces

<code>college</code>	activities, campus life
<code>columbia</code>	space shuttle and the STS program
<code>comic5</code>	the funnies
<code>consumers</code>	interests, product reviews, complaints
<code>cooks</code>	recipes
<code>crypt</code>	data en/decryption
<code>cse</code>	computer science education
<code>cycle</code>	motorcycles and related products and laws
<code>dcom</code>	data communications hardware/software
<code>decus</code>	DEC Users' Society
<code>emacs</code>	EMACS editors
<code>eunice</code>	SKI Eunice system
<code>flame</code>	flaming (complaining) on any topic
<code>followup</code>	followups to articles in <code>net.general</code>
<code>games</code>	games and computer games
<code>emp</code>	the Empire political/economic simulation
<code>frp</code>	fantasy role playing
<code>go</code>	the strategy game Go
<code>hack</code>	the Hack computer game
<code>pbm</code>	play-by-mail games
<code>rogue</code>	the Rogue adventure game
<code>trivia</code>	trivial knowledge
<code>video</code>	video games
<code>garden</code>	gardening
<code>general</code>	important and timely announcements of interest to all
<code>graphics</code>	computer graphics, art, and animation
<code>ham-radio</code>	amateur radio practices, contests, events, rules
<code>info-terms</code>	terminals
<code>invest</code>	investments of money
<code>jobs</code>	job announcements, requests
<code>jokes</code>	
<code>d</code>	discussions on the content of <code>net.jokes</code>
<code>tias</code>	children, their behavior and activities
<code>lan</code>	local-area network hardware/software
<code>lang</code>	computer languages
<code>ada, apl, C, f77 (FORTRAN-77), forth, lisp, mod2 (Modula-2), pascal, st80 (Smalltalk-80)</code>	
<code>legal</code>	legalities and the ethics of law
<code>lsi</code>	large-scale integrated circuits
<code>magazine</code>	proposed new mail/network standards
<code>mail</code>	gatewayed from the ARPA header-people list
<code>headers</code>	gatewayed from the ARPA MsgGroup list
<code>msggroup</code>	

Continued

math	discussions and puzzles	5c1	general science
stat	statistics	5f-lovers	science fiction lovers' newsgroup
symbolic	symbolic algebra	singles	single people and their activities
med	medicine products and regulations	social	like net.singles, hut for everyone
micro	microcomputers of all kinds	sources	posting software packages and documentation
16K	National Semiconductor 32000 series chip	bug5	bug fixes and features
432	Intel 432s	games	recreational software
6809		mac	Macintosh software
68K		space	space programs and research
apple		sport	spectator sports
atari		startrek	Star Trek, the TV show and the movies
chm	Commodore Business Computers	std	IEEE, ANSI standards
cpm	CP/M operating system	suicide	suicide, laws, ethics, and its causes and effects
hp	Hewlett-Packard	taxes	tax laws and advice
mac	the Apple Macintosh and Lisa	test	for testing network software
pc	IBM personal computers	text	text/word processing
ti	Texas Instruments	theater	performers and audiences of the theater
trs-80	Tandy TRS-80	travel	travelling all over the world
zx		tv	
misc	discussions too short-lived for other groups	drwho	Dr. Who
motss	homosexuality (Members Of The Same Sex)	soaps	soap operas
movie5	reviews and discussions of movies	unix	Unix neophytes group
5w	Star Wars saga(s)	unix-wizards	Unix hug reports and fixes
music	music lovers' group	usenix	Usenix Association events and announcements
classical		usoft	universal (public domain) software packages
folk		veg	vegetarians
synth	synthesizers and computer music	video	video and video components
net-people	info about people on the net	wanted	requests for things that are needed
news	discussions of USENET itself	source5	requests for software, termcap entnes
adm	comments directed to news administrators	wines	wines and spirits
b	discussion about version-B news software	wobegon	The Prairie Home Companion radio show
config	system down times and interruptions	women	women's rights
group	discussions and lists of newsgroups	only	postings by women only (read by all)
newsite	new site announcements	works	workstations
sa	comments directed to system administrators	fa	from the ARPAnet
stargate	satellite transmission of news	arms-d	arms discussion digest
nlang	natural languages, cultures, heritages	arpa-hhoard	ARPAnet bulletin board
notes	notesfile software from the University of Illinois	human-net5	computer-aided communications digest
origin5	evolution versus creationism	info-kermit	Kermit communications software
periphs	peripheral devices	info-mac	Apple Macintosh micros
pets	household animals	info-term5	all sorts of terminals
philosophy		info-vax	DEC's VAX line of computers
physics		info-vlsi	Very Large Scale Integrated circuits
poems		laser-lovers	laser printers, hardware and software
politics		poli-5c1	politics and/versus science
theory	theory of politics and political systems	railroad	real and model train fans' newsgroup
puzzle	puzzles, problems, and quizzes	tcp-ip	TCP and IP network protocols
railroad	real and model train fans' newsgroup	telecom	telecommunications digest
rec	recreational/participant sports/hobbies	mod	moderated groups
birds	budwatching	map	net maps/routing announcements/software
boat	boating	news	maps of the Usenet network of news sites
bridge	bridge card game	uucp	maps from the uucp mapping project
coins	coin collecting	motss	newsgroup on gay issues and topics
disc	flying disc (frisbee) activities	movies	reviews and discussion of movies
nude	naturist/nudist activities	music	reviews and discussion of things musical
photo	photography	newslists	news-related statistics and lists
scuba	scuba diving	singles	articles of interest to singles
ski		sources	postings of public domain sources
skydive		test	testing of moderated newsgroups
wood	woodworking	std	standards
religion	religious, ethical, and moral implications of actions/christian, jewish	c	C language standards
research	research and computer research	mumps	standards for MUMPS
roots	genealogy	unix	Unix system features and hugs
rumor	postings of rumors		

interesting patterns in the requests. In particular, I noticed that I got many requests from different people at the same company. In one case, I received 10 requests from people at one company not known for work on the topic. This, I determined, was a market signal that could be used to determine the interests, and perhaps even the plans, of research and development at the company. The rapid publication of abstracts and gathering of requests using the network, compared to paper publication, can thus provide extremely sensitive information.

I should point out that it was never my intention to use the network for corporate espionage. It was simply the unexpected interest of so many people at certain companies in specific subjects that made it clear to me that networks could be used for such purposes. None of this strategic information was ever made available to anyone other than myself, and this is the first public statement of the methods. No doubt, it could have some effect on company policies on the use of networks.

NET SURVEYS

Netnews is an efficient way to gather and disseminate information, but I have also found it useful for gathering opinions. I once created a netnews group called `att.ch1` for AT&T employees interested in computer-human interaction. One use I found for this group was to distribute design alternatives and collect opinions about which was best. In two days, while I was free to work on other tasks, I received 14 replies, many of which were detailed and useful. Compare this method to more traditional methods for getting feedback—I could have spent my time walking around for a few days to get people's opinions, but I would not have contacted many people far from my location or people I did not know. I could have used internal publication, but my ideas were not

well formed at the time, and in any case, internal publication can take many months. The ease of requesting information, here combined with the ease of responding to requests, made netnews an unusually efficient research tool.

GROUP PROGRAM CONSTRUCTION

Another positive use of the net is group construction of programs. In late 1984, I posted a C program version of a shell script originally posted

Collaborative efforts from different environments makes for rapid program development.

to the news group `net.sources` in 1982. The conversion of the shell script to C was to add functionality and improve performance. In a short time, people at other sites were posting versions based on mine. These had bug fixes, portability enhancements, and some presented new features. I incorporated these changes in my version of the program and posted a request that people send changes to me so that the same program would not have separate development tracks. To encourage this, I used a time-tested

Netnews is primarily a non-profit medium with expenses shared by new subscribers, not news posters. Because of this, commercial advertisements have always been met with criticism, often hostile. People on the net like to read announcements of new products, but only if they are posted by uninterested parties to the appropriate news group (never `net.general`, more like `net.math.stat`).

strategy: I acknowledged the work of others in the source code and in the documentation. This seemed to work, and I received more enhancements and thanks for their acknowledgement.

The collaborative efforts of many people working in different environments made for rapid program development. First, features that worked on my local version of the Unix system and not on others were corrected quickly. Also, the addition of subtle features that come about after months of use came about in a few days because of the large user population.

Group programming could also work inside companies with trade-secret interests because news group distribution can be limited. I am not sure whether you could apply the same method to other cooperative efforts, such as writing papers; collaboration on papers might be best left to electronic mail.

The idea of posting free software to the public has to be foreign to many people. It must also be one reason why site managers maintain their netnews connections. Still, not all software is met with public approval. Sometimes, perhaps more often than not, the programs posted have bugs, and fixes have to be posted. Note that if it takes much time to fix buggy program postings, then the software cannot really be considered free. But even when programs are bug-free, some people on the net will complain, proving that whenever people choose a public profile, they stand a risk of public and even private attack, all through the medium of netnews.

CONFERENCES

One of the early uses of netnews was to organize conferences. The Usenix conference (Unix system users) has for years used the news-group `net.usenix` to post calls for papers and schedules of talks. These are usually followed by ordinary

electronic mail submissions of papers. For some conferences, summaries of sessions have been posted while the conference was running. For many Unix system users, electronic transfer, either through netnews or electronic mail, is the sole means of communication before a conference.

Although I know of no examples on Usenet, netnews could be used as a medium for a distributed conference on some topic. Instead of going to a conference in one city, the participants could meet in one news group. This is what goes on in the daily discussions of topics on netnews—one person has an idea, posts it, and waits for the comments. In many cases, the resulting comments are better thought out than those from a panel session because people have more time to think. On the other hand, *anyone* can add their opinions, so the quality of contributions is highly variable. An organized "netconference" would require much more moderation than is common to netnews; uninvited speakers could not be quieted. Even overcoming some organizational problems, people might not want to participate because of the lack of re-

spectability of the medium for publication, not to mention the loss of perquisites.

CONCLUSION

Netnews is an informal, unmoderated, rapid, and widely used medium for communication. It can be an effective tool for doing research, some more valuable and more noble than

Netnews is informal,
unmoderated, rapid,
and widely used.

others. But more important, because of its size and speed, netnews can support research paradigms not possible before.

To make the best of netnews as a research tool, you must recognize its human component. The contributors to the net do so not only for fame and fortune but also to help out other people. It is important for people to be acknowledged for contributions to an informal network, just as it is important in more formal forms of publication. Similarly, because it is so easy to post

items of dissent, you must show restraint when reacting to the well-meaning actions of others. You can find further reading on network etiquette in the widely circulated article nicknamed "etiquette" (Schwarz 1982). □

ACKNOWLEDGMENTS

I thank Rudy Bazelmans, Glen Bruns, and Susan Gerhart, for their comments on earlier drafts of this article.

REFERENCES

- Norman, D. A. "The Trouble With UNIX." *Datamation* 27: 139+.
- Perlman, G. "UNIX STAT: Compact Data Analysis Programs." *The American Statistician* 38:2,155.
- Schwarz, Jerry. "Usenet Etiquette: Emily Post for Usenet." Article<642@eagle.UUCP> posted to netnews group **net.general** 19Nov82 21:14:55 GMT.
- Summers-Horton, K., and M. Horton. *Status of the UUCP Project*. Presentation at the Winter Usenix Conference. Dallas, TX. Published by Usenix Associations, El Cerrito, CA.
- The Wang Institute is an accredited, nonprofit, academic institution founded in 1979 by Dr. An Wang (founder of Wang Laboratories). The Wang Institute is not affiliated with Wang Labs.
- More information about public-domain statistical software that I developed can be obtained from me at The Wang Institute, Tyng Road, Tyngsboro, MA 01879. Electronic mail can be sent via UUCP to **wanginst!perlman** (connected to **decvax**) or to **sdcs!perlman** (connected to **sdcvax**), or via CSNET to **perlman@wanginst**. Phone users can try voice communication by touching 617/649-9731.

Gary Perlman is an assistant professor of information technology. Born in Montreal, he went to college at the University of Rochester and graduate school at the University of California, San Diego (UCSD). At UCSD, he began using the Unix system and started work on user interfaces, both of which he has continued at Bell Laboratories in Murray Hill, N.J., and at the Wang Institute.

UUCP AND USENET EXPLAINED

uucp is a set of programs for communication (mostly mail and netnews file transfers) between Unix systems. The **uucp** network is a set of several thousand Unix system-based machines that transfer information using phone lines and hard wired connections. All that is needed for a machine to transfer to another is: *one*, the system name; *two*, the times to call (time of day or days of the week); *three*, the device used for the call (phone or wire); *four*, the line speed (e.g., 1200 or 300 baud); *five*, the phone number if item three is a phone, and, *six*, login informa-

tion (the **uucp** name and password). Although *low tech*, **uucp** is popular because systems can communicate with minimal hardware investment.

Usenet is a subset of the **uucp** network in which the path between any two machines is unique. This is done to simplify the routing of articles posted to netnews, and to insure that only one copy of an article is sent to a machine. I leave it to the reader to think about the possible distribution problems if machines tried to distribute the same news to hundreds of machines.