news@ftp

The FTP Software® Newsletter

Volume II, Number 2 October 1989

Inside…

Item	Page
New Product News	1
Masthead	2
Internetworking News	2
Technical Support	3

PC/TCP® Version 2.04 Released

By Janet Marcisak

Our PC/TCP Version 2 software will be going to 2.04 this month, and that means bug fixes, and new features. Among the new features for 2.04 are a new BOOTP client, interrupt 14 support for Telnet and Rlogin, multiple connection capability for **rloginvt**, an enhanced **ping** program with many new options, and an updated **host** program.

We are introducing a BOOTP client in this release. This is a bootstrap protocol facility which allows a workstation to have important configuration parameters, including its IP address, set over the network. By modifying a single file on a machine running a BOOTP server, the network administrator can now change the configuration of other machines — including gateways and servers — without changing information at every node. Our implementation is fully compliant with RFCs 951 and 1084.

Our **rloginvt** program now supports multiple connections. This means that you can have as many as 10 simultaneous connections from your IBM® Personal Computer® to remote machines. The initial connection is numbered connection 0. You can switch to other connections by simply typing the PC/TCP escape character, followed by a number from 1 through 9. This makes maintaining sessions on remote machines as easy as toggling between screens. We've also added a built-in FTP server to **rloginvt** which will allow your PC to act as a server while you are logged in to a remote system. Also, **rloginvt** now does VT220 terminal emulation in addition to VT100 and VT52. These features were previously only available in our implementation of Telnet.

Our **ping** program, traditionally used to isolate trouble on a network, now supports ICMP error handling, and all IP options, including Source Route and Record Route. These additions greatly enhance the functionality of the utility. For example, you can use the Type of Service and Record Route options for a good look at the best routes between you and any other point. The new version of **ping** also provides the ability to dump sent and received packets in hex.

The updated host program allows the DNS (Domain Name System — see RFC 1034) server to be specified, and can send most DNS query types. Users can now query Domain servers for Address, Canonical name, Host INFO, and Mail exchange records, as well as others.

Also new in version 2.04 is interrupt 14 support for **tnglass** and **rlogingl**. This support allows you to use a wide variety of commercially available terminal emulators for DEC, HP and Tektronix® terminals. We have also added 132-column support to **supdup,** and enhanced **tnvt** to allow the user to set the

continued

news@ftp

Volume II, Number 2 ¥ October 1989

Published by FTP Software, Inc.
Editor: Nancy Connor.
Communications pertaining to this newsletter should be directed to:

news@ftp
FTP Software, Inc.
26 Princess Street
Wakefield, MA 01880-3004
Phone: (617) 246-0900
FAX: (617) 246-0901
Telex: 981970
Answerback: FTP SFTWRE UD
Electronic mail: news@ftp.com

FTP Software welcomes suggestions and contributions from readers. Submissions for the next issue of the newsletter must be received no later than December 15, 1989.

© 1988, 1989 by FTP Software, Inc. Permission to use, copy, modify, and distribute this publication for any purpose and without fee is hereby granted, provided that this copyright and permission notice appear on all copies, the name of FTP Software, Inc. not be used in advertising or publicity pertaining to distribution of the material without specific prior permission, and notice be given that copying and distribution is by permission of FTP Software, Inc.

continued from page 1

end-of-line characters and to support PrtSc and Ctrl-PrtSc for screen dumps. We ve also implemented a umask , which the user can set via the configuration file, for use by both the **rcp** program and InterDrive™.

Fixes in the 2.04 release include improvements to **smtp**, **smtpsrv**, **ftpsrv**, **mail**, **tar**, **install**, the Development Kit, NetBIOS, and the PC/TCP kernel. The kernel has been enhanced to improve ARP handling, reliability and speed, and now supports the February, 1989 version of IP Security options.

These new features and fixes will be available as an upgrade to anyone with the PC/TCP or PC/TCP Plus package. Contact our sales department for further details.

Protocol Futures

by John Romkey

My home computer understands a lot of different protocols. TCP/IP, which also means ICMP, UDP, ARP, Telnet, FTP, SMTP. UUCP, which gives me some of the functions of TCP/IP, on a different network substrate. RS232. ST506. It might understand SCSI or ESDI.

Of course, this list is a little unfair it s grouping bananas, oranges and grapefruit. TCP, RS232 and the disk controller protocols all do different sorts of things. Protocols like TCP, UUCP and OSI do pretty similar things, even though they do them in different ways.

Over the last few years, we ve learned a lot about TCP and IP. Network congestion and TCP timers have been two continual nuisances in the TCP world, and Phil Karn, John Nagle and Van Jacobsen have all discovered good algorithms for dealing with these problems. We ve learned how to broadcast, how to build a distributed host name database, and presumably how to route packets. The technology for building the basic TCP/IP protocol stack is pretty well understood now, although still not quite commonplace. Much of what we ve learned carries over to other protocol stacks, such as OSI. We ll see in the next few years if anybody takes advantage of this fact.

The main reason for mentioning ST506, SCSI and ESDI is that, basically, computer users don t care which controller interface their computer has. Users have a nice (or in some cases, not so nice, but you get the point) file system that sits on top of the controller, and the only way they can tell the difference between the controllers is through price, reliability and performance.

Within a few, maybe five, years, network protocols should work in a similar way. Users don t really care how the data gets there. They want to be able to move files, login, send mail, and access their database. If the data is sent over TCP or TP4, it doesn t make much

difference to them, just as long as everything they have can communicate. To do that, the whole network will either have to speak the same protocol or will have to take advantage of translating gateways which allow similar applications protocols to run on top of dissimilar transport protocols (TCP, TP4, UDP, whatever) and still interoperate.

There s a lot of work being done in this area already, although it s not often discussed this way. There are mail gateways between the TCP/IP SMTP protocol and proprietary networking software like 3Com™ 3+™, Novell NetWare™ and Banyan VINES®, as well as between SMTP and the OSI X.400 protocol. Marshall Rose s ISODE software package allows OSI applications to run on top of TCP streams, and software has been developed using it to gateway between OSI networks and TCP networks and allow the same applications to run. Every day millions of pieces of electronic mail and network news flow between the USENET, where they re transmitted via UUCP, and the Internet, where they re transmitted with SMTP.

The network market is growing at a phenomenal rate. As long as it is possible for unsophisticated users to continue to use the technology (and this is one of the important problems TCP/IP is facing), it should not stop growing until every computer in the world is connected to a world-wide network. But this network won t be speaking one single protocol. There are strong political reasons today for this not to happen; some organizations feel they must or must not use certain protocols. OSI is a certainty in the future, if for no reasons other than political ones, but when it finally arrives (in perhaps three years), there will be a huge installed base of TCP networks, and users will be faced with installing new OSI systems that are more expensive, less reliable (because they are less mature products) and more difficult to use (because users are accustomed to the old TCP products and these will work differently). As long as they can run the new applications they want, they will keep their TCP-based networks when politically feasible and install only those OSI products they must. The networks will be made a coherent whole by the use of translating gateways.

The computer world in 1995 will look quite a bit different from the one now, with much more sophisticated applications that don t require users to know anything about the underlying network technologies and allow them to connect to a world-wide network consisting of TCP/IP, OSI, DecNET, and proprietary protocol families, all transparent to the user. Just like today, when I can run ESDI and ST506 disk drives on my home computer and never have to worry about which is which (except for the fact that one of my drives is noisier than the other).

FTP Software s Corporate Network

by Stev Knowles

FTP Software s headquarters in Wakefield is a model of complexity for most people designing computer networks to use as a shining example of how to make a network much too complicated. FTP Software maintains a complex network by design, the goal being to provide a diverse and complicated environment for the testing of our networking products. By providing this complex network as the production environment all employees use, FTP Software s products can be much more thoroughly tested than they could if we tried to maintain an interesting and diverse network only for testing purposes.

Included here is a map of FTP Software s third floor offices. The map of the second floor offices is not nearly as complex, except in one room (the developers playroom), where all the current networks have drops for the developers to test new software on.

Nothing is simple in networking, and in this environment even something as simple as an Ethernet™ is not just a simple piece of wire. FTP Software runs on one logical Ethernet, but it is physically several different runs and configurations. The downstairs Ethernet runs around the second floor and up into the machine room on the third floor. There, it connects through a thin-net transceiver to a fan-out box (often referred to by the name of the DEC device of this type, called a DELNI). The fan-out box has 10 ports, connecting several machines and devices that ex-

pect to find thick Ethernet (and to connect to it via transceivers), instead of thin Ethernet, which connects to the network via T-connectors. Machines such as the VAXes and SUN connect to the fan-out box, as do two additional fan-out boxes, to provide Ethernet service for the offices adjoining the machine room.

Attaching fan-out boxes to each other in this fashion is referred to as cascading , and is not recommended for more than two levels, as we have done. The reason for this is that cascading to more levels introduces too many problems with timing for the hardware to work correctly. There is also a repeater connected to the fan-out box, connecting the upstairs and downstairs thin Ethernet segments.

The thin Ethernet on both floors is wired into the offices in the following manner. You will notice in the diagram that there are straight runs between offices, and loops on adjoining walls. We decided to install thin Ethernet in the building walls, and run local loops inside the offices. This allowed us to easily expand the thin Ethernet, without having wires dangling down walls. A thin Ethernet drop comes down inside a wall, and terminates in a wall plate with two BNC connectors. The top connection is up-stream , or closer to the machine room. If you were to put a terminator on the upper connection, you would isolate the rest of the network from everything past you on the wire. The bottom BNC connector is down-stream from you, or going away from the machine room.

When you remove the plate from the wall in one office and look in the wall, you find the plate for the other office is right across from you, and the bottom connector of one plate is attached to the top connector of the other (which is downstream depends on which plate is closer to the machine room). As you can see, the upstairs Ethernet starts at the repeater (marked Rpt on the illustration) and ends in the office closest to the stairs. From the perspective of the users, this is all one Ethernet they never see the complexity involved in the wiring or the hardware.

Some of the complexity does help the repeater will isolate the two thin Ethernets if one goes bad or is open. This allows the people on the third floor to be unaware of the fact that the people on the second floor have broken their Ethernet. More importantly though, it allows us to test software and hardware in less than ideal conditions, allowing us to (hopefully) improve the robustness of the product when dealing with less than ideal networks.

But, as you all know, all the world is not

an Ethernet. FTP Software currently has in production use several other networking topologies, and current plans call for more introductions in the near future. The router connects the IBM Token ring (4 megabit), the STARLAN (1 megabit), the ProNET-10®, and the Ethernet together. The sales department uses the 1 megabit STARLAN, and the administrative department uses the IBM Token Ring for their every day communications uses. A 10 megabit STARLAN in the playroom is connected to the thin Ethernet for use by the developers. The 10 megabit STARLAN is connected via transceiver and the 10 megabit STARLAN cards are connected to a fan-out box in the same way the Ethernet fan-out box is used.

The ProNET-10 is used by one of the VAXes, and there are drops in the offices adjacent to the machine room. FTP Software connects to the Internet through a 9600 Baud serial line to Cambridge, MA, and we also have two internal SLIP connections one into the support office, and one into a developers office. This complexity allows for constant traffic on all of these networks and for us to observe PC/TCP interacting with itself and other TCP/IP implementations across links with widely diverging speeds, from 9600 baud to 10 megabits per second. Constant work is being done on the product to increase performance across networks and between different TCP/IP implementations.

In the near future, we will be installing a 16 megabit IBM Token ring, along with an IBM Ring Bridge. We are looking to add a system to one of the networks other than the Ethernet (we are considering an IBM RT on either the 4 or 16 megabit rings), as well as purchasing some other machines which have networking code which is divergent from the standard code being shipped from Berkeley for their version of UNIX. We are also adding drops for the IBM Token Rings and STARLAN in the support office, and are planning to add at least two more routers (one for linking internal nets, and one for a new, 56 KBit link to the Internet).

We are also running applications which have been developed outside of FTP Software to look at ways of improving the performance of the network for these applications. Not too surprisingly, things like X windows and databases use the network in a much different fashion than Telnet or FTP Software does. We currently have a Banyan VINES file server, and are in the process of bringing up Novell NetWare and 3+Open file servers for production use.

We try to keep at least one of each board that we support in production use, to allow for easier testing of new releases, and to try and pinpoint places in which the current drivers need additional work.

As you can see, we have made our environment more complex in an effort to make the product better, and to make your work easier by constantly enhancing our products to make them robust and easy to support. Things are constantly changing here, and new systems are always appearing on the networks. As FTP Software grows and offers more products and services (like 3+Open™ and NetWare compatibility), the networks here will continue to grow and become more complicated. Just be glad you don t have to deal with it.

LAN Packet Demultiplexing or How You Can Run All Those Different Protocols on One Cable

by James VanBokkelen

Most LAN administrators know quite a bit about LANs, such as how data is broken up into packets, and how each host has a hardware address by which it is known to others. Many could outline the basic mechanisms of Ethernet or Token Ring, and lay out the wiring for an office. However, if faced with the question We ve got this Brand X LAN cable all over the place, can I connect my two Brand Y machines to it? , they may well be at a loss for an answer.

As long as hardware interfaces for whatever LAN standard is in use (Ethernet, 802.5, STARLAN, etc.) are available for the new machines, and they run the same version of the same network protocol (NetWare, DECnet, TCP/IP or whatever), compatibility is usually ensured (unless one implementation has bugs). However, the reality of the situation is that this isn t always the case. One reason is that different manufacturers all too often support different LAN protocols another department might choose a different LAN O/S for their PCs, or bring IBM compatible machines into a previously all-Mac environment. Another reason is that LAN protocols are frequently specialized the original installation might be using a PC LAN O/S to share files and printers, and the new users might need to use TCP/IP over the LAN to log in to a mini-supercomputer and run linear programs or CAD software. They can t use the LAN O/S because its developers didn t include remote login in their design.

Fortunately, the designers and implementors of most of the major LAN standards had this issue in mind. LAN cabling is expensive, and LAN bandwidth is usually much larger than any single host could use, so it is a natural reaction to make the resource as widely shared as possible. The first level of sharing allows many different clients and servers using protocol A to use the same cable at the same time (multiplexing). The second, which is designed into all but a few LAN standards, allows the machines speaking protocol A to share the cable with others using protocols B and C, and even allows sophisticated machines to use several different protocol families at once.

How does this multiplexing work? When you look at how data moves over the LAN media, you find that all LAN technologies break each machine-to-machine data flow up into a series of separate messages, or packets. Each packet has a source and destination address, and framing to separate it from other packets. This allows many different hosts to share a common wire to communicate, without the overhead and cost of setting up separate circuits between each pair of hosts. It also makes it easier to design higher-level protocols that guarantee delivery of data even in the face of transmission errors. Thus, the first level of LAN multiplexing is based on packets, and individual machine addresses. The recognition of hardware addresses and framing/unframing packets is almost always done in hardware, to reduce the load on the host CPU.

To implement the second level of multiplexing requires some way of telling protocol A packets from protocol B packets. At first, this might not seem necessary, unless some of the machines want to speak several different

protocols at once, because hardware addresses would seem to ensure that packets from protocol A are only received by machines that speak A. However, almost all LAN protocols make some use of broadcast packets, sent to a special hardware address which all machines receive. This is most often done to locate machines whose hardware address isn t known, or to locate a server of some sort. Since hardware addresses are usually assigned arbitrarily, as interfaces are manufactured, and change when interfaces are replaced, broadcast protocols are very attractive compared to long tables of hexadecimal numbers on each LAN node.

There are several different second-level demultiplexing mechanisms: one of the earliest, and still the most widely used, was that selected by DEC, Intel and Xerox when they defined Ethernet. Each DIX (or Bluebook) Ethernet packet has a 16-bit Ethernet packet type field in the header immediately after the destination and source addresses. All XNS packets have an Ethertype value of 0x600, IP is assigned Ethertype 0x800, DECNet Maintenance & Operations Protocol uses 0x6002, and so forth. Every DIX Ethernet protocol, from widely used public specifications like IP to private protocols used by a single product from a single manufacturer, must have a unique number assigned to it, so it is a good thing that there are 65,000 available. Ethertypes were initially allocated by Xerox, but the IEEE has taken over that function[1].

The second most widely-used demultiplexing mechanism is the IEEE packet header format specified in 802.2. 802.2 headers are in use on 802.3 (IEEE Ethernet), 802.4 (Token Bus), 802.5 (IBM-type Token Ring) and FDDI networks. In each case the 802.2 header immediately follows the physical-layer header, which varies according to the media in use. The 802.2 header uses the same amount of space (16 bits) for demultiplexing information as DIX Ethernet, but the IEEE chose to divide it up differently. The first byte is the Destination Service Access Point (DSAP), and the second is the Source SAP (SSAP), where SAPs roughly correspond to Ethertypes. Because a number of bits have been reserved to define special SAPs, only 24 unique values are available for demultiplexing different standard protocol formats.

The 802.2 header is normally either 3 or 4 bytes long, depending on which IEEE Logical Link Layer Control method is in use. LLC1 uses 3 bytes, and provides a simple datagram service. LLC2 uses 4, and provides reliable link-level delivery of messages. The LLC2 scheme is attractive to designers of small-scale LAN protocols, because it can greatly simplify a lightweight transport layer. However, LLC2 is much less useful in internetworks, where many different LAN and WAN segments are connected by packet routing nodes. Most implementors of large-scale protocols have chosen to use LLC1 for its lower overhead, since the presence of different media types and routers requires a separate end-to-end reliable delivery mechanism even if LLC2 is supported on all links for all paths.

Because the number of available SAPs is so small, the IEEE has been very reluctant to assign values to protocols it feels are unimportant, or restricted in scope. While several vendors have simply appropriated SAP values for their private protocols, others have made use of the Sub-Network Access Protocol within 802.2. This is identified by a reserved SSAP/DSAP value, and adds a 3 byte Organizationally Unique Identifier to the normal header. The TCP/IP community obtained an OUI of 0x000000, and defined it as being followed by a 16-bit DIX Ethertype, in order to provide access to any existing Ethernet protocol on all 802-type LAN media. The eight bytes of demultiplexing information this represents is the largest in common use today.

Other LAN media may use 802.2 headers, or type fields ranging from 8 to 32 bits, but the principle is the same; soon after a packet is received, low-level software can look at some well-defined part of the packet header to find out which protocol it is destined for. If this host has no higher-level software to deal with

1. Regrettably, neither organization has ever made the list of assigned types public, which sometimes complicates Ethernet debugging. Several different lists of known Ethertypes exist, maintained by various commercial and educational network sites, but they are necessarily incomplete.

a particular packet, it should be discarded quickly, so as not to waste CPU cycles.

Of course, as everyone who has ever used a computer knows, they are rarely perfect. In the LAN world, there are certainly careless vendors, ignorant vendors and self-centered vendors. The careless kind can be recognized when their software hangs or misbehaves when confronted with a packet it doesn t understand, or when a machine crashes and jams the LAN cable until it is powered off. The ignorant didn t fully understand the standard they implemented, and send packets that violate the specification, or ship hardware that can only talk with each other and not with other vendors cards.

Self-centered vendors are sometimes the hardest to cope with. Using an unallocated packet type or SAP value can lead to protocols that won t share a cable or an interface. Implementing part of a standard but discarding the rest as inefficient has been known to cause conforming implementations to report an error for every packet they receive. Perhaps the most common instance of self-centeredness is abuse of broadcast facilities, producing systems that don t scale up well. Frequent broadcasts are ok in a cozy little community of 20 personal computers, but not when you have 400 nodes representing 5 departments on the cable, particularly when someone suggests a per-packet charge to allocate overhead costs better .

This presents a dilemma to the LAN designer or administrator if all of the different protocols abide by the rules, you can get a lot more return out of an expensive cabling job, improve interdepartmental communication and use specific networking protocols for specific jobs. If they don t, you get headaches, and unhappy users.

My opinion is that even if you can segregate each function right now, it won t last. Somebody will need to do their next CAD project on a graphics workstation instead of a PC, but they'll want to send their finished drawings to the PC users plotter. Purchasing, Manufacturing, Engineering and Accounts Payable will want to share an inventory database. Somebody will notice the cost of the tape drive to back up each individual minicomputer. Given this, any resource that helps you design, purchase, install and maintain large, heterogeneous networks is very valuable. Books, courses, conferences and seminars each have their place, but the most powerful resource is other LAN-using sites. If you can get access to one of the national news/mail/conference networks, either commercial, like CompuServe or The Source, or non-commercial like Usenet, do so. Find the LAN-related forums with high signal-to-noise ratios and post your questions, answering in turn when others want information you have. In addition, or if you can t do it electronically, check out local computer groups, go to vendors user-group meetings, ask salesmen for references and develop your own network of resources. The diversity of the LAN-using community practically ensures that someone else has experience with the hardware or software combination that you re interested in.

PC/TCP and FTP Software are registered trademarks, and InterDrive is a trademark of FTP Software, Inc.

VINES is a registered trademark of Banyan Systems, Inc.

IBM and IBM Personal Computer are registered trademarks of International Business Machines Corporation.

NetWare is a trademark of Novell, Inc.

ProNET-10 is a registered trademark of Proteon, Inc

3Com, 3+ and 3+Open are trademarks of 3Com Corporation.

Ethernet is a trademark of Xerox Corporation.