Title

THE EFFECT OF DISTRIBUTED COMPUTING TECHNOLOGY ON WIDE AREA NETWORK CAPACITY REQUIREMENTS

Author

Hall, D.

Publication Date

1987-02-01

Lawrence Berkeley Laboratory

UNIVERSITY OF CALIFORNIA, BERKELEY

Information and Computing Sciences Division

Presented at the Federal Coordinating Council
on Science, Engineering and Technology,
San Diego, CA, February 17–19, 1987

THE EFFECT OF DISTRIBUTED COMPUTING TECHNOLOGY ON WIDE AREA NETWORK CAPACITY REQUIREMENTS

D. Hall, W. Johnston, M. Hutchinson,
M. Rosenblum, and D. Robertson

February 1987

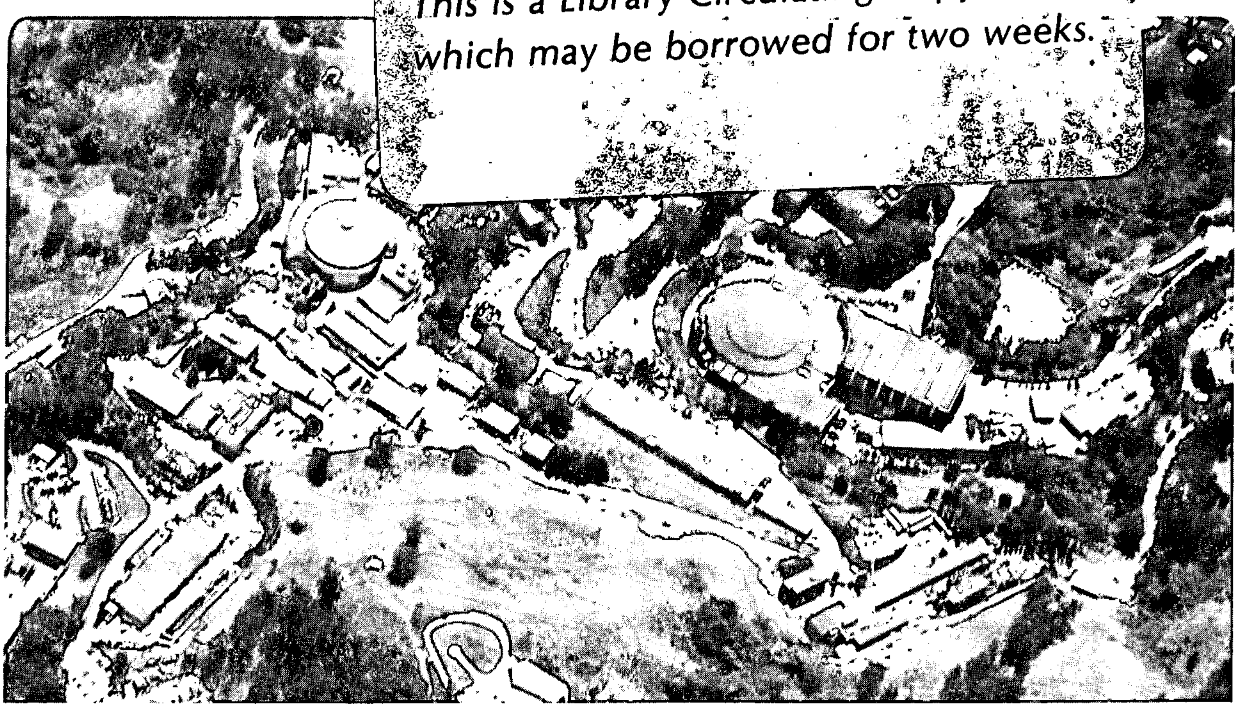

Prepared for the U.S. Department of Energy under Contract DE-AC03-76SF00098

LBL-22948

The Effect of Distributed Computing Technology on Wide Area Network Capacity Requirements

Dennis Hall, William Johnston, Marge Hutchinson, Mendel Rosenblum, David Robertson

Advanced Development Projects
Information and Computing Sciences Division
Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720

February 1987

This Research was supported by the Energy Science Advanced Computation Program of the Office of Energy Research under contract DE-AC03-76SF00098

Abstract:

This report identifies a need to increase wide area network capacity by as much as three orders of magnitude over the next ten years. These increases are necessary to support new distributed computing products. Such products increase productivity, but are currently available only on local area networks. There is no technical reason for limiting these products to tightly constrained geographical areas, however. They can operate perfectly well over any terrestrial distance provided sufficient bandwidth is available. Such bandwidth is available today with fiber optics. To quantify capacity requirements, network traffic generated by this newer technology is compared with traditional traffic in a local network environment. An extrapolation to wide area networks is made. Speculation about the long term future of distributed computing technology and its effect on network capacity requirements is offered. It is argued that an increase of network capacity by one order of magnitude is sufficient to accommodate new distributed computing technology on existing wide area networks. Two orders of magnitude are needed to accommodate a fully integrated distributed system such as interactive graphics. Three orders of magnitude are needed to accommodate increases in hardware speed anticipated in the next five to ten years. Availability of a highly integrated, nationwide distributed computing service would significantly increase the competitive edge of the United States in science and computing.

Background:

Picture a scientist using a modern, high-performance workstation. One workstation window is opened to a supercomputer located fifty miles away. The supercomputer has been programmed to compute trajectories for a beam of heavy ions from accelerator description parameters stored in a file. The parameter file is displayed on a second workstation window.

Every few seconds, the supercomputer sends a thousand position vectors and a thousand momentum vectors to the workstation. The scientist notices that at each successive step, the momentum envelope is gradually expanding, yet the spatial envelope is holding constant. This indicates heating, and means the particles cannot be focused accurately on their target.

The scientist now opens another workstation window. An interactive program is invoked that allows the beam descriptor parameters to be adjusted while satisfying physical constraints expressed as differential equations. A new value for the field gradient in the focusing magnets is established, and the trajectory computation is restarted. The program halts immediately with the message: 'Similar values were tried previously: Run 145.' The scientist reviews the results of run 145, which have been conveniently displayed by the program, then tries an entirely new value for the magnet setting. The process restarts successfully. This time there is no indication of heating, and the simulated particles proceed to a tightly focused target area.

By using a computer model, an accelerator design flaw has been detected early, and an indication of the cure has been found. The scientist now redirects the graphic display files to a video recorder so that the results can be studied by the full design team at tomorrow's review. Members of the team, some of whom are located at remote sites, are reminded of the upcoming meeting via electronic mail. A few of the images are selected to remind the scientist of the main simulation results. These are printed on a high quality laser printer located a few steps away, and shared by members of the design team. Additional copies are sent to the remote collaborators.

The preceding scenario is possible using traditional distributed computing technology widely available today. However, these older facilities are more difficult to use and more likely to cause errors than newer products now emerging in the marketplace. These newer products currently operate only in the high speed, low latency environments of local area networks. But, their utility is by no means limited to geographically constrained environments. On the contrary, these newer services are useful over any terrestrial distance, if sufficiently high bandwidth is available. In the next section we review both traditional and more modern distributed computing products.

Distributed computing products - a review:

The following distributed computing facilities have been in common use for at least fifteen years, and could all be used in the above scenario. These are the so called *traditional* distributed services.

- *virtual terminal:* Virtual terminal facilities (e.g. [1]) provide interactive access to remote machines through a computer network. In the example, a virtual terminal facility could be used to access the remote supercomputer.
- *file transfer:* File transfer (e.g. [2]) allows files to be sent from one machine to another. In the example, file transfer could be used to send the particle positions and momentum vectors to the video recorder. It could also be used to send the parameter file from the workstation to the supercomputer and vice versa.
- *remote job entry:* Remote job entry (e.g. [3]) allows batch jobs to be submitted to remote machines. In the example, remote job entry could be used instead of virtual terminal for controlling the remote supercomputer.
- *electronic mail:* Electronic mail (e.g. [4]) is the computer analog of ordinary mail. It allows the exchange of electronic "letters." In the example, it could be used to remind the members of the design team of the forthcoming design review.

While the above services are extremely useful, most of them create an unnecessary interface layer between the user and the remote resource. The paradigm of these older services is to provide a remote service through a visible network access mechanism. These were designed in the days of low performance networks, when network bandwidth was scarce. By making the access mechanism visible, users were made aware of the resource they were consuming. This of course creates extra work, provides opportunities for mistakes, and causes a certain amount of frustration. It has the undesirable effect of reducing productivity.

High speed local networks have abundant network bandwidth. As a result, more fully integrated distributed services have emerged. In these services, the network is invisible. The following are a sample:

- *distributed printing:* The emergence of high speed local area networks along with low cost laser printers has dramatically changed the way printing gets done (e.g. [5]). It is now economically feasible to allocate high quality printers to a relatively small group, and to locate these printers in the user's work area. By attaching these to a network, output can be routed to printers thousands of miles away or a few feet away. Users merely request printing service, the routing is invisible.
- *network file system:* A network file system (e.g. [6]) makes files available uniformly throughout the network. The machine on which a file resides has no special status. In our example, a network file system could be used as an alternative to shipping the parameter file back and forth between the workstation and the supercomputer. Not only is the time and trouble of shipping files back and forth saved, but more importantly, errors that arise from inadvertently using an outdated copy of the file are eliminated. Users no longer need to be aware of where their files reside. They merely access files in the usual way, location is invisible.
- *remote procedure call:* A remote procedure call (e.g. [7]) is just the ability to call a procedure (or subroutine) on a remote machine. In the above scenario rpc's might have been used to split the trajectory computation between the workstation and the supercomputer. For example, the program running on the workstation might call several compute intensive subroutines on the supercomputer, which in turn might call graphics subroutines running back on the workstation.
- *distributed window systems:* Windows provide a point for interaction between user and machine. They increase productivity by allowing users to perform tasks in parallel. A

distributed window system (e.g. [8]) permits user level programs to perform complex graphical displays on another machine's window system efficiently. It decouples details such as scaling and positioning from the generating program. Graphical data may easily saturate even the highest capacity networks. Distributed window systems attack this problem by using a high level graphics description language (e.g. [13]) for communication. This reduces network bandwidth requirements while increasing graphical display functionality. The effect is to make available high speed, high quality graphics on machines (supercomputers for example) whose graphical support system is rudimentary compared to modern workstations. While this technology is still in its infancy, it already promises a major breakthrough in the way scientists interact with supercomputers.

The thesis of this paper is that demand for these newer distributed computing services on wide area networks will increase over the next five years. In the next section we discuss our reasons for this conclusion.

Why users need local area network technology on wide area networks:

Modern, distributed computing products increase productivity. Users accustomed to these products in their local environment will want them in their extended network environment. To illustrate this we compare a traditional implementation of our supercomputer scenario with one using modern distributed computing facilities. We emphasize that both approaches are fully implementable today using off the shelf technology. We assume our scientist is using a modern workstation in either case.

In a traditional network environment, our scientist might begin by opening a graphics device emulation window on the workstation. This window would be used just like an ordinary graphics terminal attached directly to the supercomputer. Within the graphics window, a virtual terminal utility would be invoked to access the remote supercomputer. The user would "login" to the supercomputer by providing ID and password information. The remote computation would then be started by providing the names of the program and its data files in a syntax acceptable to the remote supercomputer. The number of steps required to do this would be four: open graphics window, invoke virtual terminal utility, login to supercomputer, start program.

In a modern distributed computing facility, our scientist would open an ordinary workstation window, not a special graphics window, and would immediately start the remote computation by providing the names of the program and its data files in a syntax native to the workstation. No special graphics setup, no login procedure or other conscious access to the network would be necessary. Moreover, a single syntactic framework, that of the workstation, would be used throughout. Number of steps, two: open window, start program.

Note that no login step is required in the newer environment. This is normal in local area network environment where all users work for the same institution and security is not considered a problem. For remote supercomputers, security might well be a problem. To handle this the supercomputer could request a password whenever the time between accesses exceeds a threshold. The threshold could be chosen to allow users to work unhindered so long as a reasonable degree of interaction is going on. In the above scenario, our scientist would be prompted for a password at most once.

Returning to our comparison, when the scientist notices the beam is heating, the supercomputer computation is stopped in both cases. The next step is to adjust the parameter file. To adjust the parameter file in a traditional environment, our scientist first opens a new workstation window. The parameter file is retrieved from the supercomputer by invoking a file transfer utility. This of course requires providing a login sequence (ID and password) to the supercomputer. The parameter adjustment program is then invoked. Once a new value for the field gradient is established, the scientist returns the file to the supercomputer. The file transfer utility must be reinvoked and the login sequence must be repeated before the file can be returned to the supercomputer. Number of steps, five: login, retrieve file, adjust parameters, login, return file.

To adjust the parameter file in a modern environment, our scientist simply invokes the parameter adjustment program in the same window used to run the supercomputer computation. Number of steps, one: adjust parameters.

If the parameter file is large, and if the network is typical of traditional wide area networks (i.e. 56 kbit/sec land lines), transfer time could increase frustration. However, this effect is independent of the utilities used. It is an argument in favor of high speed wide area networks regardless of the sophistication of facilities.

The next step is to restart the supercomputer computation. Except for syntax, this step is the same in both environments. The program halts immediately with the message: "Similar values were tried previously: Run 145." At this point, the scientist using the traditional environment feels the first real pangs of frustration. Four of the previous five steps must be repeated. (The file needn't be retrieved from the supercomputer initially because it hasn't been changed.) The scientist in the modern environment only repeats one step, the actual parameter adjustment.

This time the program completes successfully. To create the video display in a traditional environment, the scientist must first retrieve the output from the supercomputer then direct it (as a local file) to the video recorder (two steps). In a modern environment the scientist simply directs the output file to the video recorder (one step). Sending mail to the design team, and selecting frames for printing are done the same way in both environments (two steps). However, in a traditional environment, sending copies to remote collaborators would require electronic mail or file transfer rather than a simple print command.

The scientist using the traditional network environment has performed eighteen steps while the scientist using the modern environment has performed seven. The effect on productivity is obvious. Run setup time is reduced or eliminated because all resources (remote and local) are accessed uniformly. Errors are less frequent because there are fewer opportunities for their occurrence and because a single command syntax is used. Frustration levels are lower because less time is spent waiting for results, and because low level tasks such as shuttling files back and forth between machines have been automated. As scientists become accustomed to these modern facilities in their local environments, demand for similar facilities in wide area networks will increase.

The modern distributed computing environment we have described above is in effect a single, integrated, nationwide "supercomputer." It would be accessed uniformly from anywhere on the network. Its total power would be enormous. Such a facility, available to the national scientific community, would create a technological and scientific environment superior to that of any country in the world. It would help to maintain the nation's competitive advantage in computing for decades to come. In the next section we quantify the effects of creating such a distributed computing environment on wide area networks.

A comparison of traditional and modern distributed computing traffic

The central theme of this paper is that significant increases in network capacity are needed if local area network technology is to be extended to wide area networks. Our experience in adding such facilities to the LBL local area network is outlined in the appendix. We feel this experience provides a forecast of what might occur if modern distributed computing services were extended to wide area networks.

We observe that network file system traffic per host is about an order of magnitude higher than traditional traffic. This is based on observations of diskless workstation traffic. Diskless workstations represent file traffic that can be expected in wide area networks when users must access files from more than one machine (as in the case of our scenario). Therefore, a wide area network that operates comfortably at 56 kbits/second might need a megabit per second (i.e. a T1 channel) to support a network file system or other modern protocols that function at this level of the operating system.

We further observe that more highly integrated services such as that represented by Sun's memory swapping protocol for diskless workstations *(network disk)*, create an order of magnitude higher load than the network file system. Although network disk would not be used on wide area networks because the cost of network bandwidth is much higher than the cost of local disks, it provides a tightly coupled service at a deep level of the operating system. As such it forecasts future distributed computing traffic on local area networks. We conclude that wide area networks operating at T1 speeds might need a 10 megabit fiber optic link to support traffic from future highly integrated distributed computing services with performance characteristics similar to Sun's network disk protocol. High volume interactive graphics between a supercomputer and a workstation is an example.

So far our analysis has been based on performance of existing distributed computing facilities on existing hardware. The future will certainly bring increases in hardware speeds as well as more highly integrated network software. We think it is reasonable to project a factor of two increase every three years in available cpu power for the next ten years. Therefore, the wide area network load can be expected to increase another order of magnitude in ten years just from faster hardware.

In all, we project an increase of three orders of magnitude in wide area network capacity requirements. In other words, we think that in ten years scientists could use 100 megabit links from coast to coast to access a vast array of national scientific computers as a single, integrated "supercomputer." This would significantly increase the competitive edge of the United States in science and computing.

Conclusions:

(1) Productivity can be significantly increased by extending modern distributed computing facilities to existing wide area networks.

(2) As these facilities become commonplace in local area networks, demand for equivalent services in wide area networks will develop.

(3) To accommodate today's network file systems and other highly integrated distributed computing products, a factor of ten increase in network capacity is needed.

(4) To add software products anticipated for two to five years from now such as high volume interactive graphics, an increase of another order of magnitude is projected.

(5) To assimilate hardware speeds expected in five to ten years, an increase of yet another order of magnitude is forecast.

(6) In all, an increase of three orders of magnitude in wide area network capacity requirements are projected for the next ten years.

(7) Availability of a highly integrated, nationwide distributed computing service would significantly increase the competitive edge of the United States in science and computing.

Appendix:

Experience with distributed computing traffic at LBL

The LBL local area network consists of a single logical ethernet spanning about half the physical area of the site (several square kilometers). To isolate and minimize traffic, the ethernet is physically divided into six segments. These are joined by bridging devices with address filtering. This confines network traffic with sources and destinations on the same segment, to that segment. A single probe, therefore, can only see network traffic on one such segment. This survey is limited to statistics on two of these segments: CSRLAN, the Computer Science Research segment, and CSLAN, the Computing Services segment. It is further restricted to TCP/IP traffic only. DECNET traffic, the dominant traffic on these networks, is not examined because DECNET currently provides only traditional distributed computing services on ethernets.

The two parameters used to characterize network load are packet rate and data rate. Tables 1 through 4 summarize the load on CSLAN and CSRLAN in time intervals ranging from 21 hours to 52 hours. Network traffic is summarized for all the traditional protocols as described in the preceding section. We have included distributed printing in the traditional services because today's wide area networks carry printing traffic, although often it is disguised as file transfer traffic. The left side of the tables characterize traffic as seen by the network. The right side characterizes traffic as seen by an "average" host.

Two lines in the tables are not described in the preceding section. The miscellaneous category covers 18 relatively uninteresting protocols ranging from the *internet control message protocol* to the *time* protocol for synchronizing host clocks. The user protocol collects a variety of user developed protocols. Some use Sun's remote procedure call facility. We expect such use to increase as remote procedure calls become easier to use and more widely available.

Key to protocol abbreviations	
vt	virtual terminal
ft	file transfer
ml	electronic mail
rje	remote job entry
pr	distributed printing
usr	user defined protocols
misc	miscellaneous
trad	all the above (vt-misc)
nfs	network file system
nd	network disk

Table 1: Traditional traffic on CSRLAN - 51.99 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
vt	124	10.0	0.8	50.2	23.6	161	12
ft	59	0.6	0.2	2.8	7.6	19	8
ml	73	0.5	0.1	2.7	2.4	15	2
rje	18	1.4	1.0	7.1	30.0	157	108
pr	18	0.7	0.3	3.6	8.0	78	28
usr	17	4.0	0.6	20.3	17.5	475	67
misc	106	2.7	0.4	13.3	10.9	50	6
all trad	238	20.0	3.3	100.0	100.0	167	27

Table 2: Traditional traffic on CSRLAN - 27.52 hour sample							
service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
vt	93	9.2	0.7	37.0	13.0	197	14
ft	25	0.9	0.4	3.8	7.2	74	30
ml	58	0.7	0.1	2.7	1.6	23	2
rje	14	6.0	3.0	24.1	56.5	854	431
pr	17	1.0	0.3	3.9	5.7	112	35
usr	22	4.2	0.5	17.1	9.9	385	47
misc	70	2.9	0.3	11.5	6.1	81	9
all trad	161	24.8	5.3	100.0	100.0	308	66

Table 3: Traditional traffic on CSLAN - 41.99 hour sample							
service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
vt	79	8.0	0.8	30.0	14.9	203	19
ft	25	0.6	0.2	2.2	3.8	46	15
ml	54	0.4	0.1	1.5	1.0	15	1
rje	22	4.0	1.8	14.9	35.5	364	165
pr	32	9.3	1.7	34.6	34.0	581	108
usr	12	2.7	0.3	9.9	6.3	444	53
misc	66	1.8	0.2	6.8	4.4	55	6
all trad	154	26.9	5.1	100.0	100.0	348	66

Table 4: Traditional traffic on CSLAN - 20.74 hour sample							
service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
vt	38	6.8	0.5	13.9	6.1	360	28
ft	13	0.1	0.0	0.2	0.4	14	4
ml	39	0.4	0.1	0.9	0.6	22	2
rje	14	11.4	5.7	23.2	63.1	1630	807
pr	29	27.5	2.4	55.8	26.2	1895	162
usr	11	1.1	0.1	2.2	1.5	196	24
misc	51	1.9	0.2	3.8	2.2	73	7
all trad	100	49.2	9.0	100.0	100.0	984	179

Table 4 shows a significantly higher value of remote job entry and distributed printing load. These services are the two most heavily used in the traditional set, and use is increasing. Remote job entry traffic derives from the familiar remote shell [9] command in the Unix [10] environment. It allows commands to be executed on remote machines and files to be copied across the network. This latter use is functionally equivalent to file transfer except that no login is required. In the LBL environment, most remote job entry traffic is generated by disk backup daemons that wake up in the middle of the night to copy Unix disks onto the central VMS file cluster. The distributed printing service is also popular. It supports the Laboratory's distributed computing service that currently produces about a hundred thousand pages per month of printed output. The percentage columns show relative network traffic on the segment.

The packet and data rates for "average" hosts may be used to estimate the increase in network traffic that would be brought about by adding a protocol to a machine. Note that since all traffic is between two or more hosts, we double the network values to compute hourly rates per host. Clearly, remote job entry and distributed printing have the most significant effect in the traditional set.

The last line summarizes traditional network traffic. The network sees 20-50 thousand packets per hour and 3-9 megabytes of data per hour around the clock. An average host sees up to a thousand packets per hour and 30-180 kilobytes of data per hour around the clock. Note that the average traffic rates are much smaller than the rates for some individual protocols. This is because the denominator in the equation, the number of hosts, includes all machines in the sample. Thus, average traffic is highly biased toward protocols that run on the greatest number of systems. We next show the effect of adding modern distributed services to this environment.

Tables 5 through 8 show traditional network traffic together with two modern distributed services: network file system and network disk. Network file system is the network file system described above, and network disk is Sun's network disk protocol (proprietary). Network disk provides a memory swapping service for diskless workstations. It would not be used on wide area networks because the cost of network bandwidth is much higher than the cost of local disks. However, it provides a tightly coupled service at a deep level of the operating system. We expect to see more such highly integrated services emerging in local area networks over the next two to five years. In particular we feel that the network disk performance provides a preview of what might be needed for high volume interactive graphics between a workstation and a supercomputer. Systems such as Andrew [11] and Mach [12] provide a preview of other new products. Therefore, we feel network disk provides a good predictor for future services that will evolve first in local area networks and then be desired in wide area networks.

Table 5: Modern traffic on CSRLAN - 51.99 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
trad	238	20.0	3.3	30.9	8.0	167	27
nfs	19	5.0	2.7	7.8	6.6	529	285
nd	16	39.6	34.9	61.3	85.4	4952	4360
total	238	64.6	40.9	100.0	100.0	542	343

Table 6: Modern traffic on CSRLAN - 27.52 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
trad	161	24.8	5.3	43.9	16.5	308	66
nfs	18	4.7	2.8	8.3	8.5	521	307
nd	15	27.0	24.3	47.8	75.0	3598	3244
total	161	56.5	32.4	100.0	100.0	702	403

Table 7: Modern traffic on CSLAN - 41.99 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
trad	154	26.9	5.1	68.1	32.5	348	66
nfs	8	1.0	0.4	2.5	2.5	245	99
nd	9	11.6	10.2	29.4	64.9	2575	2271
total	154	39.4	15.7	100.0	100.0	512	204

Table 8: Modern traffic on CSLAN - 20.74 hour sample

service	hosts	kpkts/hr	mbytes/hr	pkt %	byts %	pkt/hst/hr	kbyts/hst/hr
trad	100	49.2	9.0	79.9	43.9	984	179
nfs	10	1.6	1.0	2.6	5.0	324	202
nd	8	10.8	10.4	17.5	51.2	2691	2612
total	100	61.6	20.4	100.0	100.0	1232	408

On CSRLAN network file system traffic per host is clearly about an order of magnitude higher than traditional traffic, and network disk traffic is an order of magnitude higher than network file system. On CSLAN, this doesn't show as clearly because the workstations on CSLAN are primarily used as front ends to VMS machines, and because of the previously mentioned heavy use of remote job entry and distributed printing. Nevertheless, the network disk traffic is a factor of 20 higher than all traditional traffic. We conclude that a wide area network that operates comfortably at 56 kbits/second might need a megabit per second (i.e. a T1 channel) to support a network file system. Similarly, wide area networks operating at T1 speeds might need a 10 megabit fiber optic link to support traffic from future highly integrated distributed computing services with performance characteristics similar to Sun's network disk protocol.

References:

[1] "Telnet Protocols," *The Arpanet Protocol Handbook,* Defense Communications Agency, Washington, D. C. 20305, Jan 1978, pages 89-264.

[2] "File Transfer Protocols and Standards," *ibid,* pages 265-332.

[3] "Remote Job Entry Protocols," *ibid,* pages 379-420.

[4] "Mail Protocol and Standards," *ibid,* pages 333-379.

[5] Johnston, W. E., and Hall, D. E., "Unix Based Distributed Printing in a Diverse Environment," *Summer 1986 USENIX Conference,* Atlanta, Georgia, June 9-13, 1986.

[6] NFS Protocol Specification, *Networking on the Sun,* Part No: 800-1324-03, Revision B of 17 February 1986. Sun Microsystems, Inc., 2550 Garcia Avenue, Mountain View, CA 94043.

[7] RPC Protocol Specification, *ibid.*

[8] *NeWS Preliminary Technical Overview,* Sun Microsystems, Inc., 2550 Garcia Avenue, Mountain View, CA 94043, October 1986.

[9] rsh-remote shell, *Unix Programmer's Manual,* 4.2 Berkeley Software Distribution, Virtual VAX-11 Version, March, 1984.

[10] Ritchie, D.M., and K. Thompson, "The UNIX Time Sharing System," *CACM,* Vol 17, No 7, July 1974.

[11] Morris, James H., et al, "Andrew: A Distributed Personal Computing Environment." *CACM,* Vol. 29, No. 3, March 1986.

[12] Rashid, R., "Threads of a New System," *Unix Review,* August 1986.

[13] Warnock, J. and Wyatt, D.K., "A Device Independent Graphics Imaging Model for Use with Raster Devices," *Computer Graphics,* Vol. 16, No. 3, July 1982.