SLAC – PUB – 4234
February 1987
(M)

Analysis of Network Statistics *

R. L. A. COTTRELL

Stanford Linear Accelerator Center
Stanford University, Stanford, California, 94305

Abstract

This talk discusses the types and sources of data obtainable from networks of computer systems and terminals connected by communications paths. These paths often utilize mixtures of protocols and devices (such as modems, multiplexors, switches and front-ends) from multiple vendors. The talk describes how the data can be gathered from these devices and protocol layers, consolidated, stored, and analyzed. The analysis typically includes merging information from data bases describing the network topology, components, etc. Examples of reports and displays of the information gleaned are shown, together with illustrations of how the information may be useful for trouble-shooting, performance measurement, auditing, accounting, and trend prediction.

Contents

1 Notice 3

2 Introduction 3

3 What are we Analyzing? 4

*Work supported in by the Department of Energy, contracts DE-AC03-76SF00515

Invited talk presented at the Conference on Computing in High Energy Physics, Asilomar, CA., February 2-6, 1987

4 What Data is Available? 6

5 How is the Data Gathered? 8
5.1 Audit Trail 8
5.2 Local-Recording Components 10
5.3 Probing 13
5.4 Scheduling Data Gathering 13

6 Data Bases 14
6.1 Configuration Data 14
6.2 Statistics Data 15
6.3 Analyzed Data 15

7 Analysis 16

8 What Do We Learn? 16
8.1 Media & Link Layers 17
8.1.1 Twisted Pairs 17
8.1.2 Ethernet 17
8.2 Network & Transport Layers 18
8.3 Session Layer 18
8.3.1 Simultaneous Terminal Sessions 18
8.3.2 Queuing Information 23
8.3.3 Connect Time 24
8.3.4 Audit Trail 25
8.4 Application Layer 25
8.5 Miscellaneous 27
8.5.1 Total Component Count 27
8.5.2 Trouble Tickets 28
8.5.3 Availability 28
8.5.4 Response Time 28

9 Distribution of Information 30
9.1 Hard Copy 30
9.2 Users Terminals 31
9.3 Mail & Interactive Messages 31
9.4 Broadcasts 31

9.4.1 Monitors 31
9.4.2 Welcome Messages 32
9.4.3 Status Messages 32
9.4.4 Pre-emptive Messages 34

10 Concluding Remarks **34**

11 Glossary **34**

12 Acknowledgements **36**

A Appendix: Analysis of BITNET Traffic at ASILOMAR **36**

1 Notice

2 Introduction

I have interpreted the title of this invited talk with a grain of salt and extended the talk to cover what statistics are available, how to gather them, how they are saved and analyzed, and what we can learn from them. I will try to illustrate with examples from real life.

Despite what a famous English prime minister once said about statistics:

> "There are three kinds of lies: lies, damn lies and statistics".
> *Benjamin Disraeli (1804-1881)*

I will try to show that there is useful information gleanable from the analysis of network statistics.

Reasons for gathering, analyzing, and displaying network statistics include:

- To provide a basis for billing or allocation.
- To understand how the network is used, i.e. by whom, how much, for what, and when.
- Identifying problems and permit their resolution before they seriously affect the system.

  In the long term this means being prepared to add or reallocate capacity to the network to remove bottle-necks or provide extra service. This has to be done in a cost-effective manner one does not want users sitting idle due to lack of access to resources, and one does not want expensive unused equipment sitting around. In other words the added capacity has to be available "just in time".

  In the short term one is looking for warning indicators of problems or soft errors, such as high error rates, that indicate areas that need attention.
- Identifying intermittent problems, such as a front end having crashed, a non-working modem, and providing automatic detection of critical problems (alerts).

3 What are we Analyzing?

For the purposes of this talk I shall define a network as "a collection of computer systems connected by communications paths". In the current case computer systems are taken in a broad sense to mean everything from mainframes to dumb terminals. Figure 1 is an illustration of a real-life network showing some of the major components involved. As can be seen, the network has a large number of components, most including at least a micro-processor with some level of intelligence. Many of these can produce statistics concerning their individual performance. In the case of the network shown in Figure 1 the components owned by SLAC come from over twenty different vendors. This large number illustrates the fact that typical

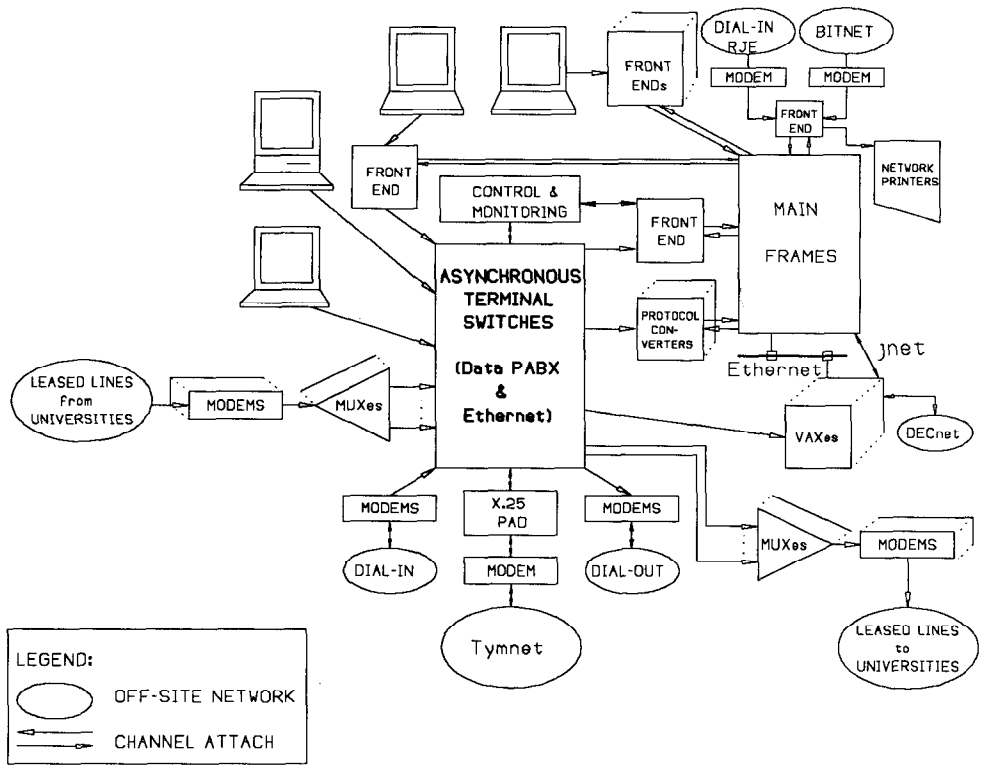

Figure 1: *The SLAC Data Communications Network. More details on the terminal side of the network can be found elsewhere[1].*

networks are not created overnight but evolve into a potpourri of equipment from many vendors. Even within a single vendor for a single component type (e.g. terminals) there can be multiple models and vintages which do not behave identically. Some of the components are not even owned by SLAC. For example, the multiplexers and modems connecting the off-site universities are owned by the universities, the off-site links are leased from the phone companies and/or satellite companies, and off-site networks such as Tymnet, BITNET and DECnet each include thousands of components over which we have essentially no control. In addition many components such as modems, front-ends, and multiplexers implement several levels of some network protocol. At the same time the complete network is probably required to support many different network protocols with arcane acronyms such as BSC, HDLC, X.25, DDCMP and protocol families such as those used by DECnet, the DARPA suite, and XNS.

Currently, there is no single standard as to whether or how each of these devices format or report their statistics. They do not even all use the same time clock. Yet in order to isolate problems in the communications paths, the network manager must gather data from most of the components at the physical level. In addition, to be able to solve systems-level problems, the manager must gather data concerning the logical network in order to acquire insight into all layers up to the user's application. Not only is this information required for reactive trouble-shooting and isolation, but also for proactive planning.

4 What Data is Available?

In the sense of network protocol layering, the network management data should be collected from all layers as shown in Figure 2.

The data available from the lower protocol layers mainly keeps track of the number of bits and bytes flowing on the links, the individual link utilization, and errors.

As one moves up through the link to the transport layers, one learns more about the framing of the data and the checksum errors, retries, and lost frames or packets.

At the session layer, data may be generated when sessions start and

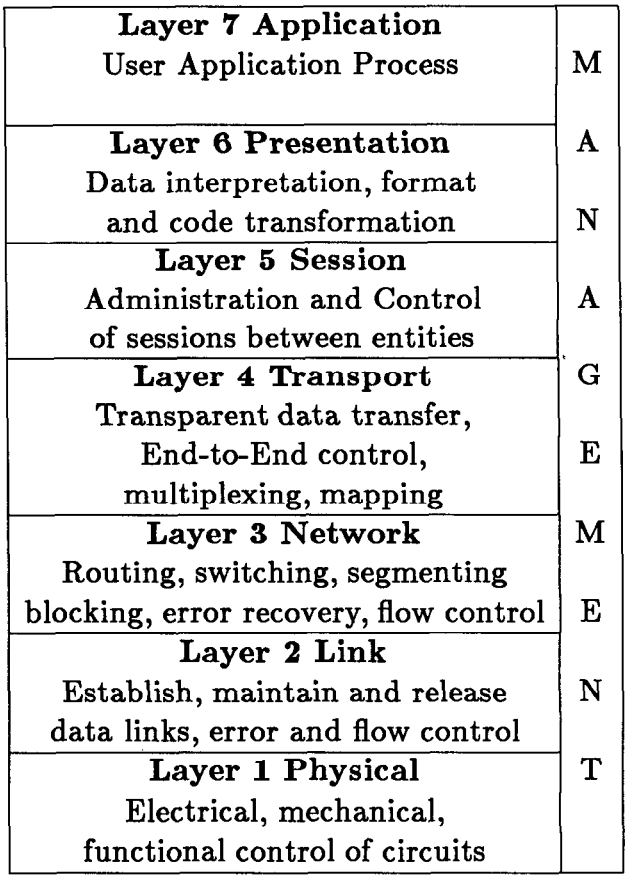

Figure 2: *The International Standards Organization (ISO) Open Systems Interconnect (OSI) Layered Architecture, showing how the network management accesses all layers.*

finish which give time stamps, type of session, services accessed, source and destination addresses involved, and how the session terminated. Sometimes this data is available as session records, i.e. one record giving the start and end times of a session. Other times there is a separate record for the start and the end of a session. There may also be records indicating queuing, and failures to create sessions due to unavailable resources or insufficient authorization.

At the application layer one is getting into the province of traditional data processing. The types of information that may be available are the volume of user data transmitted, the names of files, the print classes and printer names used, the user account names, etc.

5 How is the Data Gathered?

Figure 3 is an overview of the data flow and functions performed on the data. Many components in a network do nothing to provide data on their performance. Many simple modems and dumb terminals fall into this category. More sophisticated components may emit an audit trail of data, or collect data locally.

5.1 Audit Trail

Data switch components or Data PBX such as the Micom 600, or a network controller such as the Bridge Communications Inc. NCS/150 Network Server, can be requested to emit an audit trail. The audit trail usually consists of time-stamped records giving data on an event-by-event basis. An example of audit trail records from an NCS/150 is shown in Figure 4. The records typically contain the initiating and destination addresses, a record type identifier (e.g. session connection failed, queued, dequeued, connected, or exception). together with relevant information such as the initiating and destination addresses or error counters. A good terminal-use audit trail will contain: "session records" with the initiating and destination addresses, the session length, and a time stamp; information on the number of ports in use and available for a given group of ports or service class (sometimes referred to as hunt-group or rotary); and indications of abnormal events. Typically

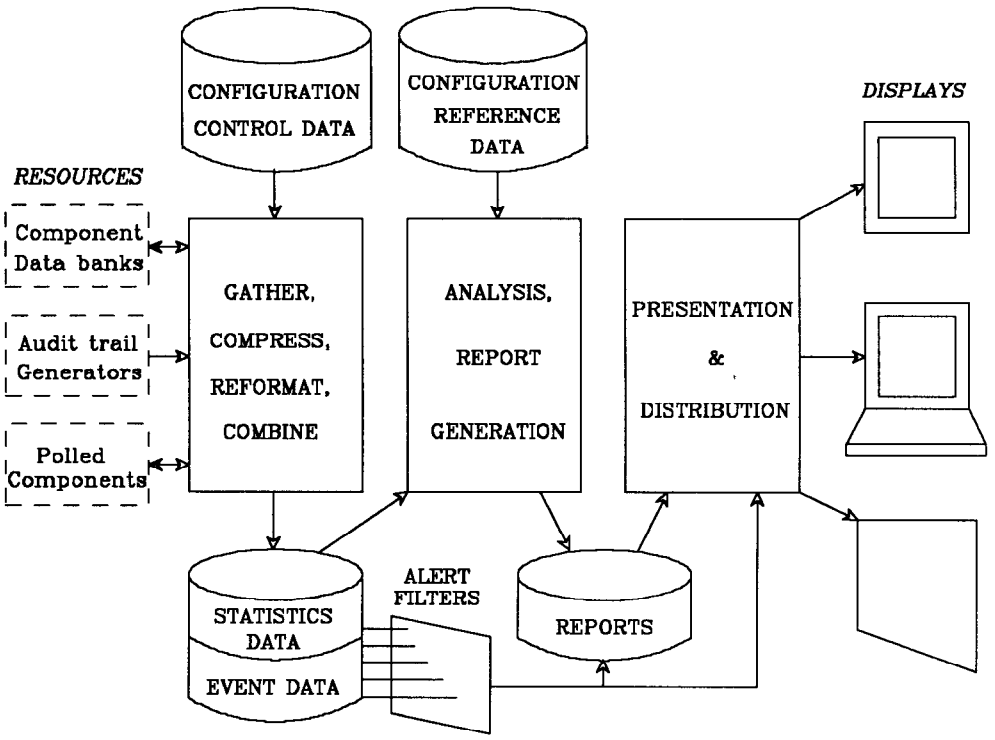

Figure 3: *The data flow for Network Statistics Analysis. On the left are the hardware and software components that generate the data (the resources). This is gathered and organized and then analyzed for easy understanding. Included in the data are high-priority items classified as alerts, which identify more serious problems. The analyzed data and alerts are distributed to the appropriate people and displayed on terminals, monitors, and hard-copy.*

they are designed to be written to a hard copy terminal or serial printer. The idea is that the network manager can later scan through the hard copy to discover problems, etc. In order to automate problem hunting, to keep a machine readable copy of the audit trail information, and provide the current status of part of the network to users and network personnel, it is necessary to connect up a computer to the audit-trail's output port. This computer then "listens" to the output port, reads the records, and can then filter and consolidate the data before recording it on a mass-store device. At SLAC we use a DEC PDP 11/60 to gather the connect and disconnect records from the Micom Data PBX and generate session records together with ports in use and available for each service class. The NCS/150 Network server gets its information from the Bridge Communications Ethernet terminal servers themselves by listening for statistics information emitted by the terminal servers each minute.

5.2 Local-Recording Components

Many components in a network gather their own statistics and keep them in local memory banks. Examples of such components include multiplexers, Packet Assembler Disassemblers (PADs), front-ends, terminal servers, Ethernet controllers, and computers. To integrate this data into the overall network measurements, these separate data banks need to be copied to a central repository. In the case of SLAC the central repository is an IBM mainframe. How one copies the data depends on the individual components.

1. Many components, such as Timeplex multiplexers, provide a separate RS-232 port that allows one to connect a terminal, for configuring the component and to interrogate the stored data. To copy this data to the central repository one can use an agent such as a personal computer (PC) to emulate a terminal and connect it to the component's port. The agent issues a script of relevant commands to interrogate the statistics data and records the results in a file. Later the agent connects to the central repository and uploads the file to the repository. This may be automated so that no manual intervention is required. The agent may also be by passed: the computer managing

Date	Time	Init	T	Dest	C	Data
04/12	10:47:44	&...	CF	&...	BU	
04/12	10:49:50	&...	DC	&...	NR	321 2992 ..
04/12	10:51:28	&...	CD	&...		
04/12	10:54:08	&...	RO			21 0
04/12	10:58:25	&...	DC	&...	OK	407 11774 ..
04/12	11:09:59	&...	CD	&...		
04/12	11:12:19	&...	SE			23742

Figure 4: *Audit Trail produced by a Bridge Communications Inc. NCS/150[2]. The columns labelled "Init" and "Dest" have been abbreviated. They normally contain the Ethernet addresses of the initiating and destination devices. An example of an Ethernet address is &00003140%08000200A6C!019, it gives the Internet number (preceded by an &), the Ethernet station address (preceded by a %), and the port number (preceded by a !). The column labelled "T" gives the type of record (e.g. CF = Connection Failed, DC = Disconnected, RO = Rotary call with available counts of ports and in-use, SE = Excessive IO errors). The column labelled "C" gives an explanatory code (e.g. BU = Busy, NR = No Response, OK = Normal Disconnect). The "Data" column gives information such as length of session, number of bytes transmitted, and number of errors.*

the repository can emulate a terminal itself, connect directly to the component and transfer the data directly at regular intervals.

2. A variant on the above that is used by the Dynapac PAD, for example, is to allow access to the data banks from any of the regular terminal ports on the component. The access is usually protected by means of some arcane commands and a password. Since the regular ports are probably connected to the network anyway (e.g. via the Data PBX), this simplifies connecting to the port and no dedicated hardware is required. Otherwise the method of copying the data is the same as in Item 1 above.

3. PADs also usually provide X.25 access to their local data banks via an X.25 virtual circuit. Thus if the computer managing the repository has X.25 access, it can easily upload the data.

4. Other components such as those used in IBM's SNA products, DECnet, and the CS terminal servers, typically have their own protocols for accessing the data banks. Sometimes these protocols are supported by the computer managing the repository in which case copying the data is easy. Other times one has to write some software to implement the protocol and copy the data. The latter presumes the protocol is published or can be deduced by some means.

5. Data banks kept by components such as terminal front-ends and Ethernet controllers, which are directly connected to the repository computer, can be read directly into the computer with very little pain without even involving the network itself.

When gathering the data one also has to be aware of what time period the data covers. Some components such as the Micom-Interlan NI1010 Ethernet controllers automatically clear the statistics whenever they are read. Others allow the user to clear the data separately from reading it. Others automatically clear out data that is older than some time period.

5.3 Probing

It is also possible to deduce data by probing parts of the network. Examples of this include:

- Automatically logging onto a remote computer to see if the circuit to it is operational, whether the computer is up, and to measure the trivial response time.
- Putting a circuit in a loop-back state and using a Bit Error Rate Tester (BERT) to check out the circuit and measure the error rate.
- Once a minute the PDP 11/60 at SLAC sends request packets to each CS Ethernet terminal server asking for the sessions currently in progress. From the data returned the PDP 11/60 builds a picture of all the current sessions and, based on changes, generates and saves sessions records.

5.4 Scheduling Data Gathering

In addition to knowing where there is data to be gathered and how to gather it, one must decide how often to gather it.

Data that is to be used for real-time displays of the current status of the network obviously needs to be gathered every few seconds or minutes. Typical of this data is response time, component status, service class utilization, and status alerts.

Data for performance reports, e.g. link utilization in bytes, packets, or retries, probably is only needed on a daily or weekly basis.

It is also necessary to take into account the permanence of a data bank. The CS Ethernet terminal servers, for example, only keep data for a sliding window covering the last 24 hours, so it must be gathered at least daily. Components that are subject to losing their data in case of failure (such as a power spike) may need to be polled often enough so that little data is lost in case of a failure. Similar considerations apply to components that can overflow their data counters.

Finally one has to consider the cost of the polling. Polling a remote computer through a Public Packet Switched Network (PPSN) such as Tymnet, or via the phone will cost real money. Even accessing a service that you are

not charged for can mean the service and some of the network bandwidth is unavailable to others.

6 Data Bases

There are several types of data that are required to be stored by or available for the network statistics analysis system. The data probably exists in many files, some of which may be simple sequential text files, while others may be managed by sophisticated data base management systems such as SPIRES or VAX/Rdb. Some of the files may be kept on different computers in the network and copies made for analysis. Some of the data may be part of another data base system that is not oriented for the analysis of network statistics, or maintained by network personnel. For example, some or all of the component descriptions may come from a property control data base. The major types of data required are given below.

6.1 Configuration Data

This data includes:

Contacts: Contains the names, addresses, and phone numbers of contact people such as system managers, field service personnel, etc. These names may be picked up and displayed in trouble reports.

Control: Contains information about the type and frequency of data gathering, polling intervals and timers, relative costs for alternate routing strategies.

Reference: Contains information about: threshold values; software (version number); hardware (serial number, supplier); protocols; modems (speed, phone number); costs of using lines (average Tymnet character charges and connect time costs); terminal port and line attributes (address, authorization, speed, protocol).

Topology: Contains a "map" of the network. This may be in the form of x-y coordinates and/or connectivity information. The connectivity information may be down to the level of individual wires and termination blocks.

6.2 Statistics Data

This is where the data gathered from the components is stored. To conserve storage space or to facilitate later access, the data gathered from the remote data banks may be compressed (e.g. by combining separate connect and disconnect messages into complete session records), reformatted, or combined with other data from the configuration data base, for example, before being saved. This data is typically organized by network component groups, for example PAD utilization data includes the associated modem, leased line, and probably several computer/terminal ports.

6.3 Analyzed Data

This data is the output from the analysis programs. It is typically in the form of tables and graphical information. It is used to create reports for users and network personnel. It is probably organized by group of network components and by history (e.g. by month, by day, last 30 days, year to date, last 24 hours). The kinds of data available include:

Alerts: Used to pin-point problems, identify where performance is below acceptable levels.

Performance detail: Provides information on response times, message and byte counts, errors.

Terminal Utilization: For individual terminal lines and logical groups of terminal lines and ports (service classes) to provide a history of: connect time and number of connects for terminals and terminal service classes; maximum utilization; queuing activity; failed disconnects (due to user error, authorization, time outs, etc.).

Summaries: Typically these are graphical data showing long term trends.

The volumes of data in some of these files may be quite large. For example, at SLAC:

- Each month about 60K session records are generated from the Micom Data PBX, and these typically occupy about 6Mbytes/month.

- The configuration data base describing the terminal line and computer ports holds about 4Mbytes, and the cabling connectivity data base has about 8Mbytes.
- The analyzed network tabular reports data base has about 4Mbytes of data.

7 Analysis

The magnitude of the analysis task can vary from interactively applying simple editor macros to small files, through applying a set of rules to a data base manager to select a subset of data, to large batch jobs running commercial analysis packages such as Statistical Analysis System (SAS) available from SAS Institute, Inc.

The basic techniques used in the analysis are usually conceptually simple. Typically the data is filtered by comparing with thresholds from the configuration data base; sorted, based on variables such as the originator and destination addresses, error types, accounting group, length of session, service time intervals (e.g. work hours, by month, by day), etc.; and normalized with other data (e.g. with respect to time for rates, error rates compared to throughput, useful data compared to overhead). Often the data is summarized in terms of sums, averages, standard deviations and extrema such as maximum utilization and unused facilities.

The output of the analysis may be viewed immediately or saved on disk or tape for further manipulation and display or for auditing purposes.

8 What Do We Learn?

Rather than try and vaguely indicate what one generically learns from the analysis of network statistics, I will use examples in various areas of networking to illustrate how we may use the results.

8.1 Media & Link Layers

8.1.1 Twisted Pairs

The traffic measurements (i.e. average bits/second) can be used to identify links that are close to saturation or that are hardly used and available for re-allocation.

Error measurements, if available, can indicate transmission media problems due to such things as noisy phone lines, terminals connected beyond their rated maximum distances (especially at 19.2kbps and above), faulty components in μwave links, or improperly terminated lines. Some modems can provide measurements of analog phone line status such as harmonic distortion, frequency shifts, and phase jitter. Sometimes one can use such information to switch to another line, or call in experts to look in more detail at the problem. The error may also be caused by auto-baud failures resulting in the speed being deduced incorrectly, and subsequent framing errors. This may be due to faulty auto-baud algorithms, or due to user error such as entering the wrong character at the autobaud time.

8.1.2 Ethernet

At the media layer one is interested in the utilization in terms of maximum and average percentage of available bit rate in use, packet rates and packet size distribution, and interpacket arrival time distributions. This can indicate how close to saturation the network is. In addition, it is valuable to know the error rates and type of errors (alignment errors, short packets, checksum errors). Ethernet Monitors from firms such as Excelan and Hewlett-Packard allow one to measure, display, and save such information as shown in Figure 5. Increases in error rates can indicate a badly installed transceiver or a failing station. It should also be possible to save the bad packets for future examination. This may help one to deduce the transmitting station's address which in turn may help indicate the source of the problem.

At the link layer, by looking at the source and destination addresses and protocol types in the packet headers, one can find out who is using the Ethernet and how. Figure 6 shows the kind of information that is available from such an analysis. Thus one can learn what stations are transmitting on

the network, watch for new stations or stations that have not transmitted for several days, and keep track of how much has been sent and received by each station.

8.2 Network & Transport Layers

At these intermediate layers one of the major uses of the information is for trouble-shooting and performance analysis. Figure 7 shows an example of the type of information that is available. Knowing the packet rates for your environment and typical applications, such as file-transfer and virtual terminal services, can enable you to predict what would happen if the current traffic were switched over a different network.

For example, using the CS Ethernet terminal server statistics, one can find out the number of packets (i.e. keystrokes) sent from the terminal to a full screen remote echoing service, such as IBM 3270 protocol emulators, or the VMS EVE editor. It turns out that the average such user at the peak time of day[1] makes about 1.5 keystrokes per second. In turn, this generates an average of 170 characters per second to and from the terminal. This puts certain limits on the bandwidth and packet handling rates required to support such services. Thus on a 56kbps X.25 link (such as proposed by HEPnet), one can only support about 65 simultaneous such users, and the switch must be able to handle about 300 packets per second.

One can also use topology maps to graphically illustrate the current network status, indicating what node or links are in normal operation, a warning state, or in a problem state. Such maps can be valuable to the network operator in real-time and they are also great crowd pleasers. Figure 8 is an example of such a display.

8.3 Session Layer

8.3.1 Simultaneous Terminal Sessions

An excellent indicator of capacity is obtained by analyzing the maximum number of simultaneous sessions to a given service class. When this reaches

[1]These values are based on measurements of terminal activity on the SLAC VM system. The figures for the busiest time of the day have been used. Rates averaged over the busiest minute and busiest hour of the day are about 60% and 20% of these values.

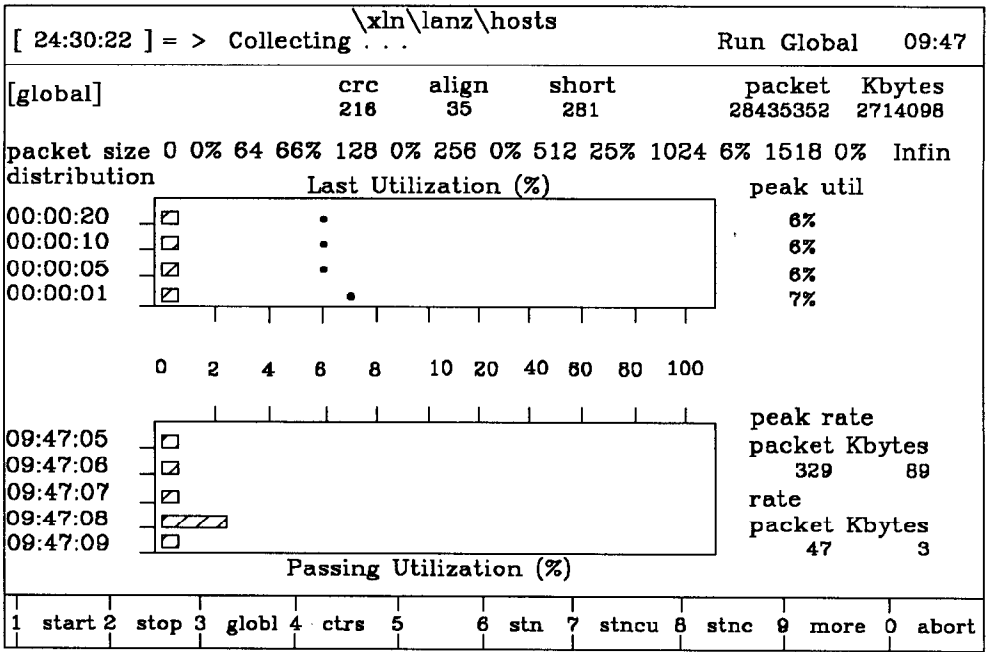

Figure 5: *Display of the Utilization of the SLAC Ethernet by an Excelan Lanalyzer[3].*

No	In	Station Address	Out	Packet		Avg-size		Errors:	
				rcv	xmt	rcv	xmt	rcv	xmt
1	1>	DECnet m'cast		350	0	122		1	
2		TPCS	2>	107	91	64	72		
3		LAN Bridge m'cast		98	0	63			
4	3>	SLD		85	70	167	102		
5		aa-00-04-00-80-a4		3	3	71	98		

Figure 6: *Part of an Excelan Lanalyzer display of the stations detected on the Ethernet, and the amount of traffic for each. The columns labelled 'In' and 'Out' give the number of packets seen in the last time interval. Commonly used station addresses are assigned alias names (e.g. DECnet m'cast); others simply display their Ethernet address.*

Time 12:51:14			
General ITP Statistics for Host: 0207 0100 27DA			
XNS ITP Protocol V01-003 Version			
0	IDP Packets Routed	0	IDP Bytes Routed
903	IDP Packets Received	5583	IDP Bytes Received
055	IDP Packets Trans.	5324	IDP Bytes Transmitted
2	Error Packets Sent	0	Packets Disc. by Router
0	Echo Responses Rcvd.	0	Router Responses Sent
41	Router Requests Rcvd.	0	Router Responses Rcvd.
401	SPP Send Requests	2504	SPP Receive Requests
192	Packet Exchange Req.	36	Packet Exchange Resp.
0	Routing Table Entries	0	Free Bytes Remaining

Figure 7: *Display of SLACnet Network and Transport layer statistics by the Micom-Interlan NETMAN package. The protocol suite is the XNS Internet Transport Protocols (ITP). The network layer protocol is the Internet Datagram Protocol (IDP). At the transport layer the Sequenced Packet Protocol (SPP), Packet Exchange, Error, Echo, and Routing Interchange Protocols are supported.*

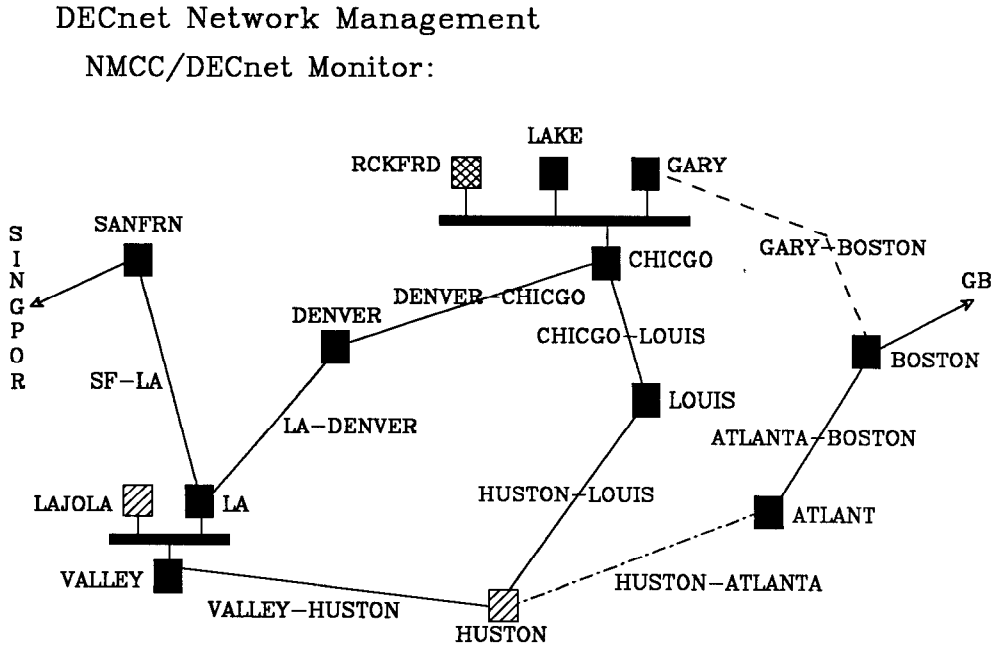

Figure 8: *A network Map of part of DEC's network produced by the NMCC/DECnet Monitor package[4]. Solid boxes and lines indicate normal operation, diagonal-line shaded boxes and dashed lines indicate a warning state, cross-hatched boxes and dot-dashed lines indicate a problem state.*

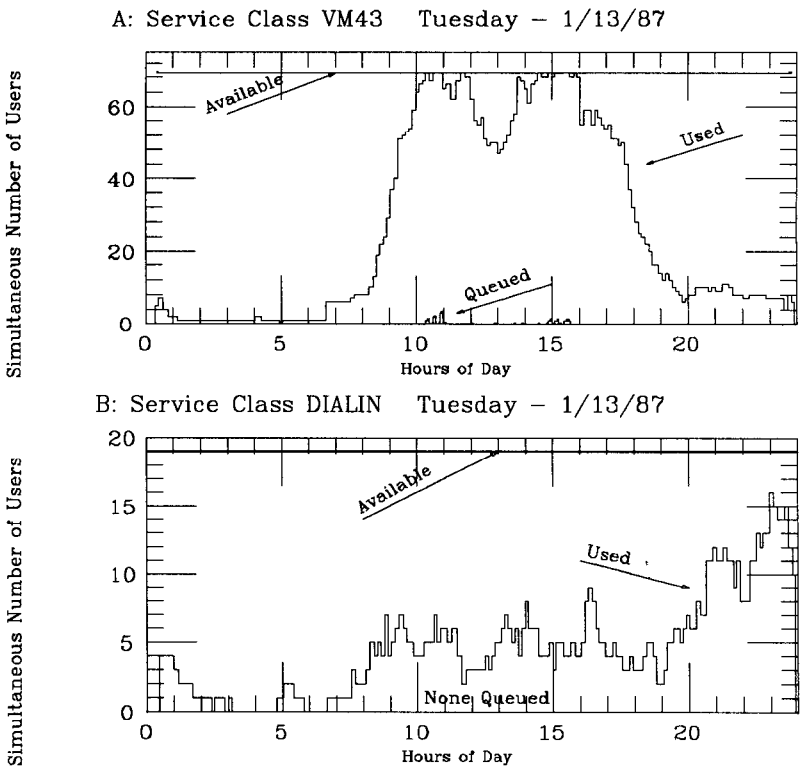

Figure 9: *Number of simultaneous sessions as a function of time of day. A: SLAC's IBM 3270 emulation services, B: Dial-in 1200bps service.*

the total number of ports available to that service class then users cannot be connected and either give up or get queued for the service. Figure 9 shows this happening for one of SLAC's service classes in mid-morning. Figure 9 also enables one to quantify patterns of use by time of day. One can easily identify in Figure 9 when most users come to work and logon, when they go to lunch and return, when they go home, have dinner, and then dial-in and logon to check a few things before going to bed. With such knowledge it may be possible to shift resources during the day to match changing needs.

Longer term trends can be detected by plotting the maximum simultaneous sessions/service class on a month by month basis. Figure 10 clearly shows the shift of users from line-by-line services to full-screen services over

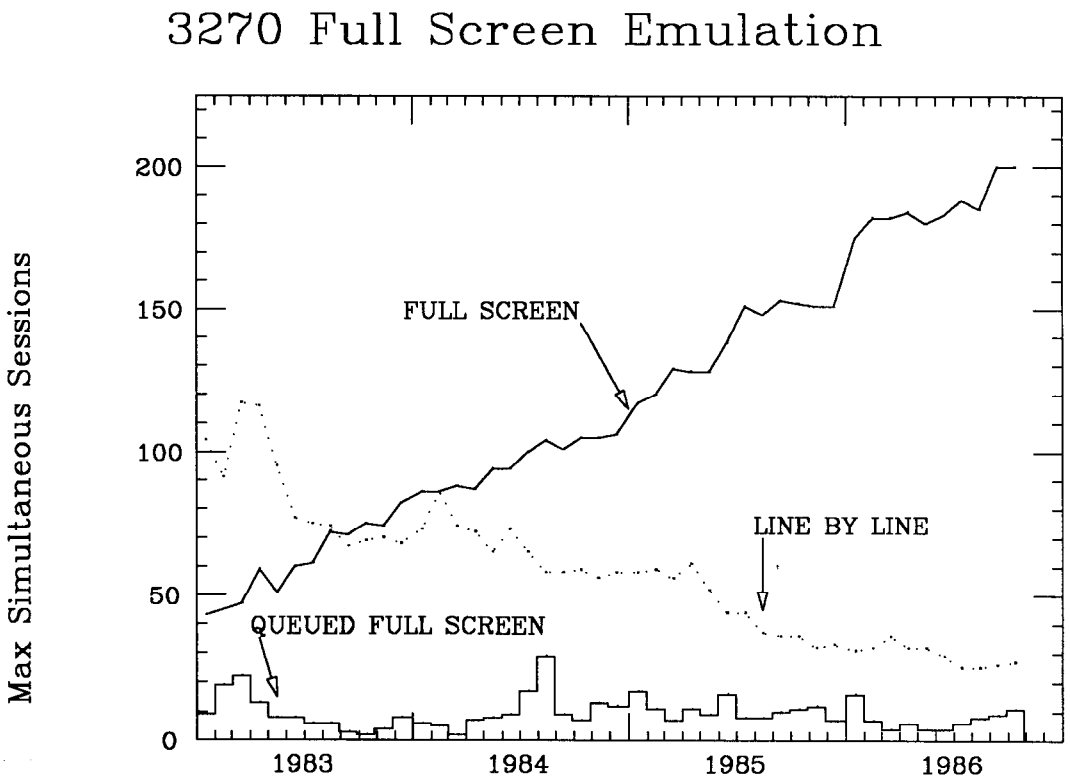

Figure 10: *Usage of SLAC's VM system in IBM 3270 full-screen emulation mode compared to that in line-by-line mode. The bottom of the graph shows the number of users simultaneously queued awaiting full-screen service.*

the last four years. By extrapolating such a plot it is possible to estimate how much capacity will be needed in the coming 12 months.

8.3.2 Queuing Information

Analyzing the queuing information and failures to set up sessions due to no free resources is a very sensitive indicator of inadequate resources. It can also be used as a measure of achieving service goals. The number of failures to connect due to inadequate resources in general should be kept to less than a few per-cent. As can be seen in Figure 10 the number of users queued can build up rapidly over a couple of months. This provides a shorter

term warning of inadequate resources complementary to the extrapolation mentioned above. The dips after a peak in the queued information in Figure 10 are due to adding extra ports to the full-screen service.

8.3.3 Connect Time

The traditional way of analyzing session records is to provide total connect hours by service, user, or user group. This may be used for allocation or billing purposes. In the case of the SLAC network, a group might be the owner of a set of multiplexed terminal lines from a university. In general, at SLAC, we do not try to allocate or bill back based on network connect time or volume of use. To do so would discourage use of the network, and connect time or volume of data is only indirectly related to the cost of running our network. However, for services which cost us real money, based on connect time or volume of use, such as those of a PPSN or the phone company, it may be important to have an independent way to verify the vendor's billing figures. This may point to discrepancies (due, for example, to an improperly terminated session), and allow a different or more detailed breakdown of the traffic. With such knowledge one may decide there are more cost effective ways to provide a service. Examples of this might be to replace some dial-in use with a PPSN, or to replace some PPSN use with a leased line.

A very important way to analyze connect time is to look at abnormal cases. Very long sessions, especially over dial-up lines or via PPSNs, may indicate a failure to properly disconnect the session at some level. This may result in large bills and/or be a security exposure if someone else later gets connected to an improperly disconnected session. Very short sessions may indicate an improperly terminated line that is babbling, or a user having problems, or an intruder trying to penetrate many different services available on the network.

Services or terminal lines which are lightly used or not used at all may be having problems or may be candidates for removing or re-allocation.

8.3.4 Audit Trail

Another major use of session records is to provide an audit trail. Say the security log on one of your computers identifies a large number of password failures originating from one of the computer's ports. The first step is probably to look at the network session records, to identify the originating address of any sessions to the computer port for the relevant time period (presuming the computer and network clocks are reasonably synchronized). In a typical case the originating address might be on a PAD coming into your network from another network (e.g. Tymnet or Telenet). Then it will be necessary to pass the information to the other network's support people, so they can search their session records to further identify the source.

8.4 Application Layer

At this level one learns who is using the network and how. Thus by analyzing the SLAC VM system's log of files sent and received on BITNET one can quickly identify that:

- 25% of the 16 Mbytes/day sent by SLAC goes to Europe.
- 30% of all files are mail. The cost/mail item sent = $0.12.
- Over 50% of the users are from physics groups.
- Over 40% of VM users/month use BITNET.

Also by plotting the number of files transferred on BITNET to and from SLAC one can see the growth over the last few years as in Figure 11.

When developing new code or after installing a new release the information can be useful to pin-point problem areas. For example a large number of failed file transfers, due to one particular reason, may indicate poor documentation leading to lack of user understanding on how to use some feature.

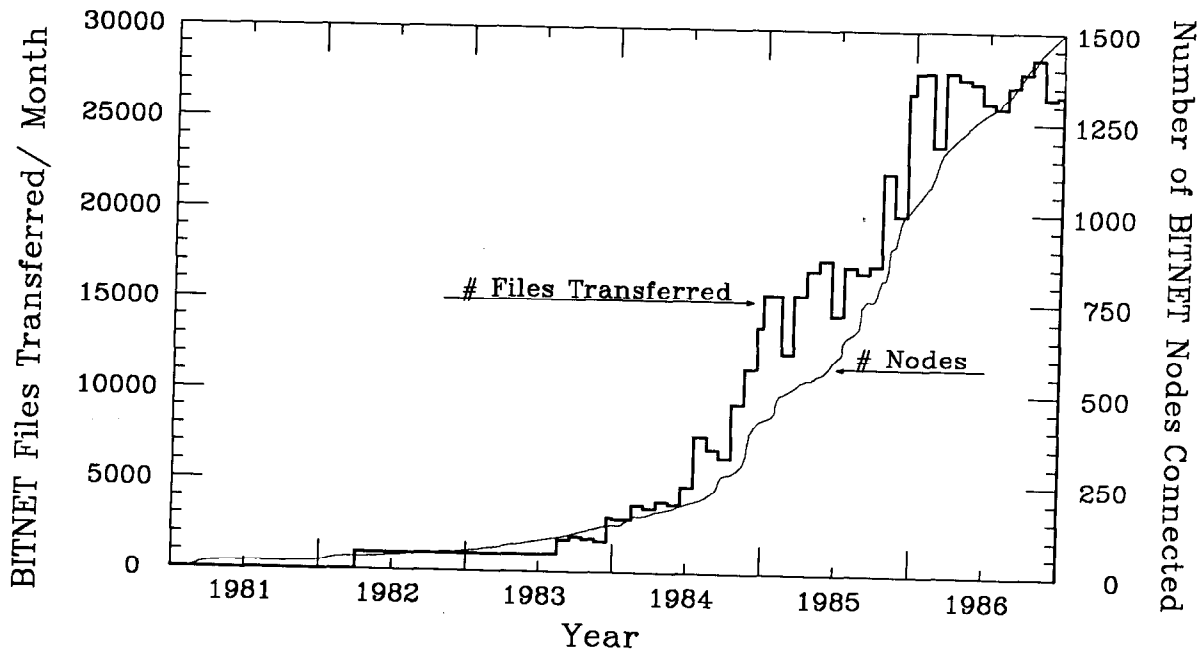

Figure 11: *The growth of SLAC's use of BITNET, and the increase in connectivity (number of accessible nodes) available.*

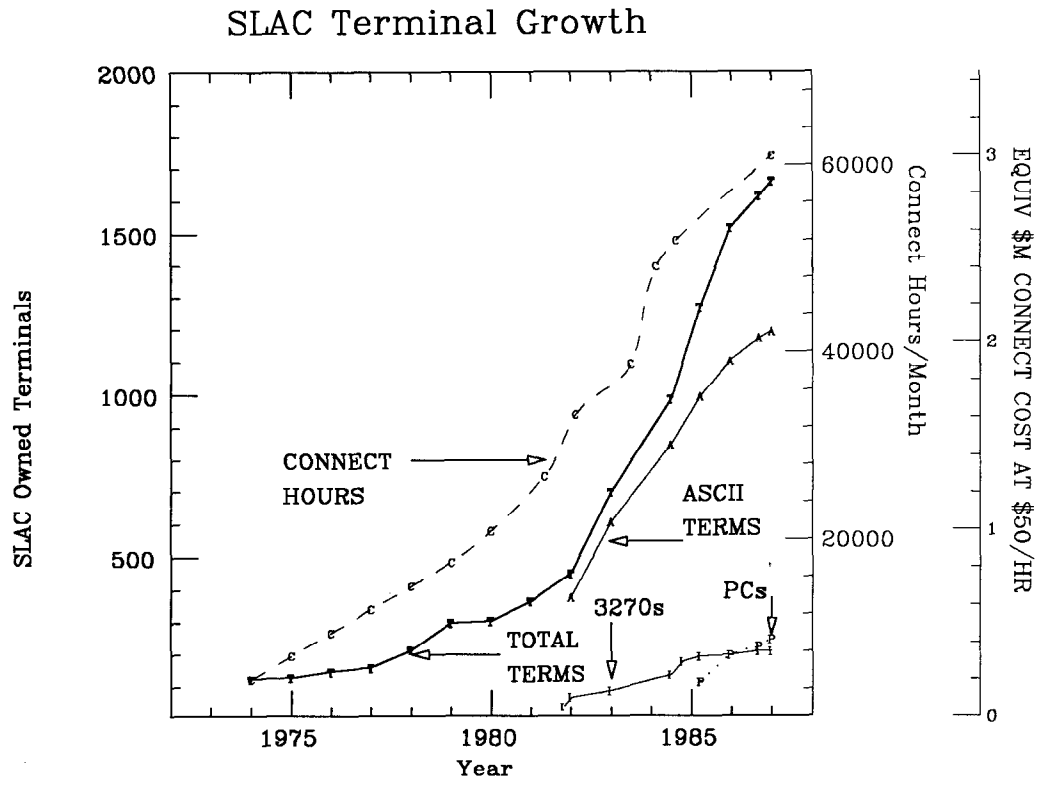

Figure 12: *Numbers of Terminals Owned by SLAC for the last decade and the connect hours per month used on the SLAC VM system.*

8.5 Miscellaneous

8.5.1 Total Component Count

Keeping track of the number of terminals purchased gives an indicator of required network capacity. This is illustrated in Figure 12 where the number of connect hours is seen to track the number of terminals at SLAC. Breaking down the numbers by generic terminal type can also suggest shifting patterns of use. Thus the growth of PCs shown in Figure 12, may indicate an increased demand to support PCs on the network.

8.5.2 Trouble Tickets

As troubles occur it is necessary to record them and then track their resolution. These records are typically referred to as trouble tickets. At SLAC we developed a SPIRES data base application to keep track of such information. Analyzing this data can reveal such things as vendor response, poor documentation or procedures, areas that need more support (e.g. ones that have a high number of open problems), change in trends, etc.

8.5.3 Availability

Availability of a network is hard to define since only some small part of the network may be down. Most of the availability problems seen by users are due either to inadequate resources (e.g. all ports for a given service class are busy), or the host is down. We keep track of queuing for busy resources and try to keep it well below 5% (i.e. the ratio of the number of users queued to the number connected for a given service class). We also keep track of the SLAC VM system's availability as shown in Figure 13.

8.5.4 Response Time

Studies [5] have shown that response time can be related to user productivity and hence to an economic value. Typically the response time seen by users is composed of computer system response time plus network response time. Measurements of the aggregate of these are important in order to determine whether additional resources are required (such as more CPU, more memory, faster transmission lines, etc.).

A more detailed type of response time is the time it takes a character entered at the keyboard of a terminal to be echoed back from a host computer and displayed on the screen of the terminal. This time becomes very important if one is using a full-screen editor, which requires the host to echo each character. If the response time increases much above 0.5 second then the feedback to the user is so slow that the user becomes confused and frustrated. Such response times are typical on PPSNs since the characters have to be packetized and depacketized and sent through many network switches on their way from and to the terminal. They can also occur if the communications links use satellites.

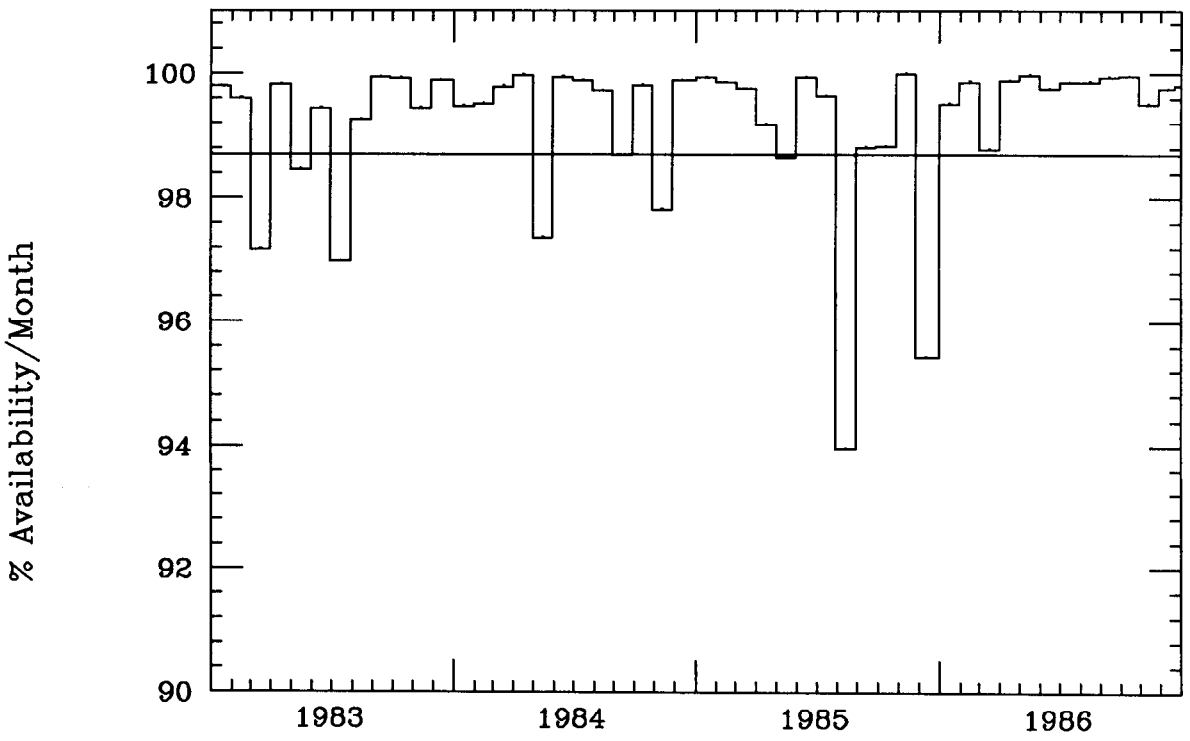

Figure 13: *The availability of the SLAC VM system for the last few years. Availability here is defined as: number of hours of unscheduled outage / total hours. The availability goal of 98.7% is shown by the flat line.*

Studies of PPSNs [6] show that for Telenet and Tymnet, these numbers extend to 730msec and 1130msec respectively. Since typical terminals automatically repeat characters (when the key is held down) at rates between 12 and 30 times/second, echo times > 33msec can result in cursor motion being bursty or erratic.

Even in local area networks considerably different responses can be observed. For example using a Black Box Dataline Monitor PC to measure the response time between a terminal and the SLAC IBM Series/1 based 3270 emulation service, we saw the time to echo a single character go from:

- 16 msec when the terminal was connected through a Micom Data PBX; to,
- 84 msec when the terminal was connected through 2400 bps Microcom AX/2400 modems; to,
- 92 msec when the terminal was connected through Bridge CS Ethernet terminal servers; to,
- 137 msec when the terminal was connected through 2400 CDS Series II modems; and to,
- 137 msec when the terminal was connected via two Micom Box Type 2 PADs.

9 Distribution of Information

There are many ways in which the reports may be delivered to the relevant people.

9.1 Hard Copy

Hard copy is often needed for reports to management, for sticking on notice boards and for embedding in documentation. We have not reached the paper-less age yet.

9.2 Users Terminals

Users and network personnel need to be able to interactively view the reports on-line at their terminal. These may be in graphical or tabular form. Typically a set of tools is used (possibly based on editor macros) to facilitate the selection of the report desired (e.g. configuration, errors, traffic), the components desired, the time span (e.g. last 24 hours, last 12 months) and the format. In addition, extensive help on what information is available and how to use the tools is necessary.

For some reports, such as alert logs, it may be possible to request further information on an event of interest. This further information may be:

- More detail on the event such as a partial dump of registers, or a longer verbal description.
- A display, from the contacts data base, of a contact person who may have more expertise in this area.
- A list of possible causes for the event and recommended actions. An example of such a display is seen in Figure 14.

9.3 Mail & Interactive Messages

When alerts are detected, the network monitoring equipment may try to warn the appropriate expert(s) such as a network operator or a system manager. If their electronic address is known (from the contacts data base), then an interactive message may be sent to the person, if logged on, or electronic mail may be sent containing details of the alert.

9.4 Broadcasts

9.4.1 Monitors

The current status of the network may be distributed to monitors or TVs strategically placed on the site. At SLAC we have a system of about 15 Commodore VIC-20s, with color TV screens, driven by 9600bps lines from a single PC. This PC acquires information about the network and feeds display information to the slaves every 2 minutes. The slaves rotate through

NCCF 3/15/83 15:26:48	
Recommended Action for Selected Event	
User Caused	Mux reinitialized during transmission
	Mux powered off during transmission
	Modem power off during transmission
Action	Correct then retry
Install Caused	Config Parm incorrect (REPLYTO)
Action	Correct Config Parm
Failure Caused	Line or mux or remote modem
	Local modem (receive side)
Actions	Run line tests
	Run remote device tests
	Run modem tests
	Contact service representative

Figure 14: *A simplified form of a Recommended Action Display provided by the IBM Network Communications Control Facility[7] of SNA.*

about 6 displays, each display staying on the screen for a few seconds. Figure 15 shows one of the displays for the network status.

9.4.2 Welcome Messages

Another way to get information to users of the network is to provide important topical information in a welcome message. Such a message is seen by the user when connecting to the network, or connecting to a given service class, or after logging onto a computer system. These messages should be terse, timely, and relevant to most users otherwise they rapidly become annoying.

9.4.3 Status Messages

For more complete information than can be put in the welcome message, longer messages can be available as a service on the network, ideally without the user having to logon to a system to display them. At SLAC, we currently have a service called STATUS on the Data PBX that allows a user to simply display the current status message at their terminal. Part of this

VMICOM
VMBRDGE
VMLINET
DIALOUT
DYNNET
DIALIN
FACSW
FLOSW
ZBRIDGE
BRIDGE
AUTO
% IN USE
Q'D
MICOM
LDC
BRIDGE
BRIDGE

Figure 15: *TV display showing a summary of the current status of the SLAC network. The horizontal bars at the top show the utilization of the major terminal service classes, the boxes at the bottom show the status of various major components of the network such as VAX nodes on SLACnet[8], the Micom Data PBX, Tymnet, etc. Blue indicates normal service, yellow is an indication of 'soft' problems, red means the component is out of service.*

information is updated automatically every 10 minutes. Such messages can also be dictated into a phone answering machine, or even converted from the ASCII text to voice, and then made available to users who call in on the phone.

9.4.4 Pre-emptive Messages

Some networks and most computer hosts allow one to send a broadcast message to all users, or a specified subset, connected to the system. Such messages must be used with care since they interrupt the users' activity and train of thought. For example, if you warn users that the system is going down in 2 minutes, and then interrupt them every 15 seconds to give them a count down, it is doubtful they will be able to clean up properly and logoff.

10 Concluding Remarks

The end goal of all this is to provide the users with the best service consistent with costs. Thus one is calling on the analyses to: improve the performance of network personnel and the productivity of users; reduce the manpower required to manage the network; and to make optimum cost-effective use of the network components.

One of the major problems today is the lack of standards, or possibly it is too many standards. It was all very well for a famous American writer to say:

> "Consistency is the hobgoblin of little minds". *Ralph Waldo Emerson (1803-1882).*

He was not a network manager, who had to find all the big (and expensive) minds required to run an inconsistent network.

11 Glossary

BSC - Binary Synchronous Communication Protocol, a system that IBM developed in the early 1960s for transmitting data over synchronous data communications facilities.

DARPA - Defense Advanced Research Projects Agency, a funding agency for the computer networking experiments performed over the "ARPANET".

DDCMP - Digital Data Communications Message Protocol. This is one of the link layer protocols utilized by DECnet.

Ethernet - A local area network system utilizing 10 MHz coaxial cable.

ISO - The International Organization for Standardization, a standards body. Among other things they have promulgated is the Open Systems Interconnect (OSI) layered model of communications protocols.

PAD - Packet Assembler/Disassembler. A functional unit that enables data terminal equipment not equipped for packet switching to access a packet-switched network.

PBX - Private Branch Exchange, a switching system that serves one company, and connects to the national telephone network. In this paper I use the term Data PBX to refer to a switching system that transmits data only, and connects terminals and computer ports together, providing port contention and selection to the terminal user.

PPSN - Public Packet Switched Network, e.g. Telenet and Tymnet in the U.S., Datapac in Canada, Datexp in Germany.

Repeater - A device used at the physical layer of the ISO layered model, that amplifies or otherwise conditions signals received from one piece of transmission medium and passes them onto another piece.

SNA - Systems Network Architecture, IBM's layered communications protocols.

RS-232 - Recommended Standard (RS) of the Electronics Industries Association for the interface between data terminal equipment (DTE) and data communications equipment (DCE) employing serial binary data interchange (August 1969).

Transceiver - A device that uses digital data to create a signal that can be transmitted over a communications media, and can receive such signals and recreate the original digital data.

XNS - Xerox Network Services, Xerox Corporation's layered data communications protocols.

X.25 - The X.25 standard defines the procedures necessary for a packet mode terminal to access the services provided by a packet-switched data network.

12 Acknowledgements

I should like to acknowledge the SLAC Network Group for producing and supporting the network monitoring tools used at SLAC. In particular this includes Teresa Downey, Tim Streater, Charles Granieri, John Halperin, and Martin Emmerson. Lois White has done most of the SAS type analysis of the consolidated data at SLAC. Ted Johnston gave a lot of inspiration for providing methods of viewing the analyzed information on-line. Billie Bennett, Lamont Barton, and Ken Martell provided the figures in this paper. I also owe a debt of gratitude to Duane Schmidt of DEC and Harry Lichtbach of IBM, who provided me with information on DECnet and SNA network monitoring tools.

A Appendix: Analysis of BITNET Traffic at ASILOMAR

Since most of the attendees to this conference were computer literate, it was decided to provide electronic mail service at the conference site. The conference organizers ordered a 9600 bps leased line in December 1986, and it was installed in time for the conference. In addition, DEC agreed to provide and install a μVAX and modems, and Joiner Associates supplied Jnet.

This Appendix describes the results of this experiment. The analysis was performed both at Asilomar (with help from Dave Shambroom of Northeastern University) and at SLAC. The μVAX was connected by the leased line to the nearest BITNET node, U.S. Naval Postgraduate School at Monterey, about 4 miles away. The μVAX was available from the afternoon of Sunday, February 1st until 10 a.m. Friday, February 6th. During

Day	Files Received	Files Transmitted	Total Files	Records Received	Records Transmitted	Total Records
Feb 1	115	97	212	9,196	1,932	11,128
Feb 2	239	224	463	11,228	5,203	16,431
Feb 3	262	303	565	9,577	9,390	18,967
Feb 4	238	325	563	16,478	13,724	30,202
Feb 5	208	292	500	5,611	11,105	16,716
Feb 6	38	36	74	1,070	2,261	3,331
Total	1,100	1,277	2,377	53,160	43,615	96,755

Table 1: *Number of files and records transmitted and received by BITNET node ASILOMAR, by day of the conference.*

this time 1,100 files were received and 1,277 were transmitted for a total of 96,755 records, or well over 5 Mbytes. Table 1 shows how these files were distributed by day of the week. Of the 205 attendees, 120 sent mail and 131 received mail. Mail was transmitted to and received from nodes throughout the U.S., Canada, Europe, and Japan. Table 2 shows the top 10 node names communicated with. On average each attendee sent or received 2 mail items/day, and the average cost per mail item sent (the line cost $358 to install and $37.54/month to lease) was about $0.30. The heaviest day for transmitting files was Wednesday and Figure 16 shows the number of files received and transmitted as a function of time of day. Coffee breaks and dinner hours are clearly visible in the upper figure.

Clearly, providing an electronic mail system at a remote conference site was highly successful. Attendees really took advantage of the system, even those who had to learn a new user interface. The only major problem encountered was that there were only four terminals available, causing long waits at prime times. The impression was that it was mainly used for meaningful communications, as opposed to "electronic post-cards (weather is beautiful...wish you were here)".

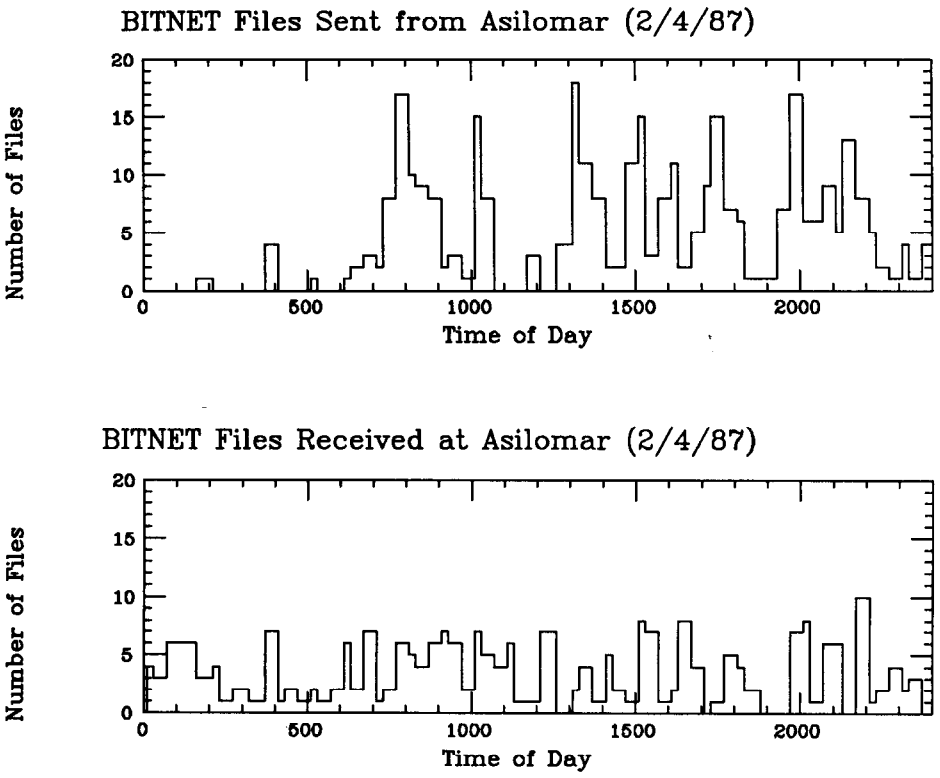

Figure 16: *Number of files transmitted and received by time of day. Note the peak transmission times before and after the sessions and at coffee breaks. The flatter distribution on the lower figure is partially due to files being received from many different time zones, including Europe and Japan.*

Node	Files	Transmitted	Received
SLACVM	437	213	224
WISCPSL	184	59	125
CERNVM	169	83	86
CERNVAX	147	94	53
SLACTWGM	124	22	102
IBOINFN	119	55	64
ANLVMS	96	75	21
JPNKEKVM	68	25	43
UKACRL	63	37	26
UCRPHYS	57	43	14

Table 2: *The top 10 nodes to which mail was sent from node ASILOMAR.*

References

[1] R. L. A. Cottrell. *A Tale of Two Networks.* Data Communications, October 1986.

[2] Bridge Communications, Inc. *Network Control Server, Installation and Operation Guide.* 09-0049-00. 1985.

[3] Excelan. *LANalyzer EX 5000E Ethernet Network Analyzer User manual.* Publication No. 4200029-00, 1986.

[4] DEC. NMCC/DECnet Monitor Introduction. AA-EW34A-TE. 1986.

[5] Walter J. Doherty and Arvind J. Thadhani. *The Economic Value of Rapid Response Time.* IBM Report.

[6] Steven M. Lauretti. *Users: See for yourselves how the public data nets perform.* Data Communications, January 1987.

[7] IBM. *Network Problem Determination Application, Version 3, General Information.* GC34-2111-0. 1983.

[8] R. L. A. Cottrell, T. Downey, H. Frese, C. Granieri, M. Huffer, L. Moss, O. Saxton, D. Wiser. *SLACnet - Implementation and Experiences.* SLAC-PUB-3894, Presented at the SHARE 66 Conference, 1986.