Columbia University

Network Integration Task Force

Final Report II

30 January 1992

TOWARDS A COMMON ADMINISTRATIVE AND ACADEMIC NETWORK AT COLUMBIA UNIVERSITY

Columbia University's Administrative Information Systems have, historically, followed IBM mainframe communications architecture. Access to online mainframe applications is mainly through 3270 terminals and emulators. The 3270 model has been implemented largely through an IBM-specific physical network of control units with dedicated point-to-point connections to the central administrative computer. This network is physically separated from the academic network, and it is dependent on closed and proprietary IBM SNA protocols. Users with dedicated SNA connections cannot access academic resources such as CLIO and ColumbiaNet.

To bring the broadest possible spectrum of administrative and academic information to the authorized user in a consistent way, the administrative mainframe can be connected to the academic network. This requires only a simple modification to the administrative mainframe: the addition of a high-speed network interface (Ethernet), installation of IBM's TCP/IP network software, and connection of the mainframe to the network. Existing applications and access methods are not affected. Users connected to the academic network will be able to access 3270 screen applications using tn3270 protocol, readily available in software for the desktop computers commonly found at Columbia: IBM PCs, Macintoshes, and UNIX workstations. Access via Rolmphone data connections through 3270 protocol converters will continue as before.

This change brings the following benefits:

- New users can be added easily by connecting them to the academic network. Special cabling, control units, and/or protocol converters are not required: only a network interface in the user's workstation which connects them to existing network wiring (such as the Rolm building wiring and the campus academic backbone on the Morningside campus).
- Old 3270-specific circuits and equipment can be retired over time, at a significant cost savings in maintenance and people-time.
- Users can access administrative and academic information in a consistent and convenient way from a single network-connected workstation, with maximum interoperability among applications.
- Connection to the academic network also allows connection to departmental host computers and to computers and services throughout the worldwide Internet.
- The move to open networking standards brings vendor independence.

Over time, Columbia's network can become simpler, more reliable, and more easily managed. Duplication of effort and redundant components will be reduced, resulting in a more effective an efficient operation. At the same time, the universe accessible from the University desktop will be expanded to an unparalleled degree.

1. Introduction

The Network Integration task force was appointed in November 1991 to examine the feasibility of integrating Columbia's administrative computing systems with the common university network. The committee members are:

Tony Cirillo, AIS, Co-leader
Alan Crosswell, AcIS
Frank da Cruz, AcIS, Co-leader
Stew Feuerstein, AIS
Brian Graham, AIS
Peter Humanik, Communication Services
Megan McCormack, Health Sciences, Co-leader
Dan Russo, Communication Services

Some of the meetings were also attended by:

Bob Bookbinder, Lamont
Paul Clayton, Health Sciences
Ken Lee, AIS
Queenie Ma, Health Sciences

1.1. Objective

To describe Columbia's existing networks and to identify the steps to move to a common, shared network and set of protocols for both administrative and academic applications, providing open, consistent, and convenient access from each desktop computer or terminal.

1.2. Executive Summary

This report examines Columbia University's current academic, administrative, and departmental computer networks and protocols and recommends a basis for integrating and managing them in a consistent way. It reflects the consensus of the task force, and was prepared mainly by Frank da Cruz, Alan Crosswell, and Stew Feuerstein.

Briefly stated, our findings are:

- A wide variety of network technologies and protocols exists at the various campuses of Columbia University.
- On the Morningside campus, the academic and administrative networks are physically separated and use different protocols and technologies.
- The academic network uses shared, common backbone cabling, rather than point-to-point dedicated connections. The academic backbone connects to the Rolm wiring, and has connections to the other Columbia campuses and their backbones and building wiring with the exception of certain buildings that are not internally wired for data.
- TCP/IP protocols are used to access academic applications but not administrative ones. SNA protocols are used to access administrative applications but not academic ones.
- TCP/IP protocols, specifically Tn3270, could be used to access most present-day administrative applications.

- TCP/IP can coexist with SNA (and other protocols) on a common network backbone.
- Open and standard protocols are preferred over closed and proprietary ones for reasons of improved support, cost savings, and vendor independence over time.
- AMS pledges, when requested by us, to support TCP/IP in its core products in a security-conscious manner.
- Many users, both on campus and at home, have no direct network access and therefore require asynchronous terminal access to central services.

Therefore, we conclude that:

- The administrative mainframe should join the common TCP/IP network using Tn3270 access to its applications.
- The physical 3270 and SNA network should be migrated, over time, to the common network subject to constraints of cost, performance, and capacity.
- Current administrative applications should be converted, over time, from 3270 screen presentation to an open and standard regime based upon TCP/IP transport.
- Future administrative applications that follow the client-server model should operate over an open networking base with no dependence upon SNA. Applications that do not follow the client-server model should be accessible from ordinary terminals, and therefore from Telnet or Tn3270.
- All applications should be designed to take advantage of the common network and of local processing power when it is available, but should still support asynchronous terminal access.
- Formal network operations and planning functions should be established, encompassing at least AcIS, AIS, and Communications Services, with an appropriate organizational and funding structure to ensure the needs and interests of the various groups, and of our users, are accounted for.

The authorized end user should be able to access administrative and academic services consistently, with equal ease and convenience, from a single terminal or workstation using a single communication method, and these services should be interoperable to the highest practical degree. We wish to promote an open, consistent electronic information environment on a reliable common network. We wish to discourage closed, proprietary solutions and unnecessary duplication of effort and resources.

Sections 2 through 5 discuss Columbia's evolving computing and communications environment. Section 6 lists our detailed recommendations. Appendix I details the options for SNA–TCP/IP coexistence and migration, and Appendix II discusses computer and network security. Appendix III defines acronyms and buzzwords. Various attachments illustrate our present and planned network configuration.

2. Columbia's Computing Environment

The administrative systems reside on an IBM ES/9121 mainframe computer, running the VM/CMS, MVS/TSO, and MVS/CICS operating systems. Administrative applications are also run on PCs and PC networks.

The central academic services reside primarily on UNIX computers: Sun, Encore, NeXT, IBM, etc, as

well as on the academic partition of the ES/9121 running the the VM/CMS operating system. Academic services are also provided on and/or accessible from PCs, Macintoshes, and UNIX workstations located in public access areas, laboratories, offices, and dormitory rooms.

The Columbia Library Information Online (CLIO) system runs under the MVS operating system on the academic partition of the ES/9121.

Various types of terminals and computers are found in the schools and departments, most commonly PCs and Macintoshes (thousands of each). Some departments have UNIX workstations (typically Sun, DEC, or NeXT). Some also have larger, shared computers, usually UNIX or VAX/VMS systems. IBM mainframes are found in a few departments.

3. Columbia's Computer Networks

The various Columbia University campuses maintain diverse computer networks running a variety of protocols.

3.1. Network Protocols

The adminstrative network uses IBM SNA and RJE protocols exclusively[1], whereas the academic network supports a rich mixture of protocols, primarily TCP/IP, IPX, Appletalk, and DECnet.

The administrative IBM mainframe host, CUVMC, communicates with the outside world via BITNET. Academic hosts use both TCP/IP and BITNET[2].

3.1.1. Administrative Network Protocols

SNA is a closed, proprietary IBM networking architecture based upon a hierarchy of mainframe computers, front ends, control units, and terminals, with provisions for certain types of program-to-program communication and management functions. It is almost universally used in IBM mainframe environments. Although SNA emulation software is available on many third party platforms, SNA is basically an architecture for mainframe communications. It is not suitable for communication among the wide variety of computers and terminals at the University.

The primary protocols used in the Administrative network are terminal-to-host protocols. The SNA PU2 and LU2 protocols support IBM 3174 controllers and IBM 3270 terminals or third party hardware and/or software that can emulate them.

Other protocols used are IND$FILE for file transfer over LU2, APPC over LU6.2 for advanced program

---

[1]IPX protocols are used within Novell networks, but SNA protocols are used to communicate between these networks and the central administrative mainframe

[2]TCP/IP is explained in section 3.1.3. BITNET is a store-and-forward networking method based upon IBM RSCS protocols, used primarily by IBM mainframes, VAX/VMS, and UNIX systems for wide-area networking over dedicated (usually leased) lines (BITNET protocols can also be run over TCP/IP). Thousands of computers all over the world are connected by BITNET. The capabilities of BITNET are limited, however, by its store-and-forward nature and the IBM record orientation of its messages.

to program communication and PU4 for communicating with other IBM SNA networks. In addition some terminal-mainframe and file transfer/printer connections use non-SNA BSC or asynchronous protocols.

IBM's NetView provides host-based network management for SNA and other IBM networks.

3.1.2. Academic Network Protocols

The underlying datalink protocol of the academic backbone network is Ethernet. Ethernet frames have a Type field that specify the network-level protocol: IP, IPX, Appletalk, DECnet, etc. This is what lets the backbone-connected routers handle multiple protocols. In the future, the backbone could be converted to FDDI or other higher-speed protocol. FDDI, like Ethernet, supports a range of higher-level protocols via a Type field.

TCP/IP is an open, nonproprietary, almost universally implemented protocol supporting a wide range of terminal-to-host and peer-to-peer functions. It forms the basis of the Internet, interconnecting approximately half a million host computers all over the world.

IPX is Novell's proprietary protocol for linking networked PCs to LAN services such as file servers, print servers, etc[3].

Appletalk is Apple's networking protocol for linking Macintoshes, Laserwriters, and file servers. It may be used on Ethernet, Token Ring, or "native" on shielded or unshielded twisted pair wiring, or it may be encapsulated within IP packets.

DECnet is a proprietary protocol used primarily by departmental DEC VAX/VMS systems.

3.1.3. The TCP/IP Protocol Suite

The TCP/IP protocol suite consists of IP (Internet Protocol) at the network (routing) level and TCP (Transmission Control Protocol) at the transport (end to end) level. Various application protocols can ride on top of TCP, including:

TELNET	Virtual terminal protocol.
TN3270	Virtual terminal protocol for 3270 emulation.
FTP	File transfer protocol.
SMTP	Simple Mail Transfer Protocol.
NFS	Network File System.
X	MIT's X Windows and related protocols.

Also, a variety of administrative and management protocols are available, notably SNMP, the Simple Network Management Protocol.

---

[3]Novell NetWare 3.11 and later can run over IP instead of IPX, and also allows encapsulation of IPX within IP and vice versa.

Applications from other protocol suites can be used over TCP transport, such as CCITT X.400 messaging and X.500 directory service, major applications in the ISO Open Systems Interconnection suite. X.500 White Pages service is currently available on Columbia's academic systems.

Finally, other protocols, including SNA, DECnet, and Appletalk, can be encapsulated within an TCP- or IP-based delivery service.

3.2. Columbia's Current Network

The current Columbia University Morningside campus data network consists of:

1. Data switching via the Rolm CBX.
2. Dedicated point-to-point circuits for CLIO terminals.
3. The administrative SNA / 3270 network.
4. The Academic backbone network.
5. Departmental LANs.
6. Connections to wide-area networks.

The networks at other Columbia campuses have different characteristics. Of primary concern to this task force is whether they have (connections to) IBM control units into Columbia's administrative mainframe, and whether they are capable of establishing TCP/IP connections to the Morningside campus. The situation is summarized in Table 3-1.

Campus	*Control Units*	*TCP/IP*
Morningside	Yes	Yes
Health Sciences	Yes	Yes
Teachers College	Yes	Yes
Barnard College	No	Yes
Lamont-Doherty	No	Yes
Nevis Laboratory	No	Yes
Harmony Hall / 2828 Broadway	Yes	No
Hogan Hall	Yes	No
Interchurch Center	Yes	No
McVickar Hall	Yes	No
Harlem Hospital	No	No
Institutional Real Estate	No	No

Table 3-1: Campus Connection Methods

Outlying buildings such as Interchurch and Hogan can be added to the TCP/IP network in various ways: Rolm 64 Kbps synchronous circuits using Rolm telephone circuits (for buildings that are on the Rolm system), leased 56 Kbps or 1.544 Mbps T1 circuits, or Columbia-owned microwave or cabling. Which method to use is a cost and performance issue. It is important to note, however, that interior building wiring is also required. Buildings whose telephone systems are provided by NYTEL (such as Health Sciences dormitories) probably do not have such wiring.

3.2.1. The Rolm CBX

The Rolm CBX allows Morningside campus users with data-equipped Rolmphones to access administrative and academic hosts via asynchronous terminal communications at speeds up to 19200 bps, to dial out to external hosts or services through a shared modem pool, and to dial in to Morningside hosts and services from outside. It also allows for synchronous connections at 64 Kbps, which can be used to make connections to the common network.

Table 3-2 shows the port configurations of the Rolm CBX as of January 8, 1991. Of the 1340 data-only ports, about 1300 are host (answer) ports, and the rest are originate ports. The LAN ports are backup circuits for the Watson, Philosophy (SIS), and Harmony Hall LANs (2 ends each).

Data-only:	1340	(857 extensions)
Student data phones:	339	(853 extensions)
Faculty/staff data phones:	1362	(854 extensions)
Inpool modems:	59	
Outpool modems:	95	
LAN ports:	6	
Total ports:	3201	

Table 3-2: Rolm CBX Data Ports

3.2.2. The Administrative Network

There are two primary means of connection in the administrative network. One is via a 3270 terminal or emulator and the other is with an asynchronous terminal. About two thirds of the network is 3270-based and the other third is asynchronous terminal-based. In both cases the Rolm wiring is used wherever possible.

The 3270 terminal architecture is roughly follows: Several (1-64) 3270 terminals or PCs with 3270 emulation cards connect via Rolm wiring to a 3174 communications controller located nearby. This controller is connected via dedicated Rolm wiring to a single NCR COMTEN front end processor (FEP) that is connected to the Administrative partition of the mainframe. There are about three hundred and fifty 3270 terminals or emulators and about forty-five 3174 controllers.

In addition there are several LANs connected to the Administrative system. The PCs on these LANs emulate 3270 terminals and communicate through the LAN to a dedicated PC emulating a 3174 communications controller via either SNA or IPX protocols. The PC emulating a 3174 then communicates with a Token Ring Gateway to the mainframe via the SNA protocol. This network architecture is an excellent intermediate step in the migration to a single integrated network, since it can use the common backbone network for communications from each LAN to the Token Ring Gateway. Two of the three LANs connected to the Administrative system utilize the current Academic campus backbone. The third LAN utilizes a leased line from NYTEL because it is located in a non-University building.

The asynchronous terminal architecture consists of asynchronous terminals (such as DEC VT320s, PCs running Kermit, etc) connected to Rolm dataphones. The dataphones utilize Rolm wiring and the Rolm CBX to communicate with Rolm Data Communications Modules (DCMs) located in the same room as the mainframe. The DCMs are connected to IBM 7171 3270 protocol converters that are connected to the Administrative partition of the mainframe. There are about two hundred Rolm dataphones capable of

communicating to the Administrative system and IBM 7171 capacity for one hundred and twenty eight simultaneous asynchronous terminal sessions.

3.2.3. The Academic Network

Simply stated, the architecture of the Morningside campus academic network is this[4]:

- On the Morningside campus, local area networks use Rolm twisted-pair wiring, jacks, and distribution frames. Rolm wiring supports Ethernet, Token Ring, and Apple Localtalk, and any higher-level protocols that run on over these datalink protocols.
- The Rolm wiring converges in the building distribution frame (BDF)[5], where LAN wires are connected to Cabletron hubs that interconnect the stations of the LAN and convert from the twisted-pair building wiring to a longer-distance medium (coax or fiber), which is run back to a router at a central location.
- On other campuses like Health Sciences and Lamont, different wiring schemes are used, but their connection to the Columbia backbone network is still through a router (or an ‘‘extended router’’ — two routers connected over long distance by microwave or T1).
- The routers are interconnected by the backbone cabling. No other devices are connected directly to the backbone.
- The routers are also connected to the worldwide Internet through our local service provider, Performance Systems International, Inc. (PSI), and gatewayed to other wide-area networks such as those based on CCITT X.25.

The hubs and routers provide the interface between the backbone network and local networks that are connected to it. These give us media independence: we can change the backbone technology without affecting the local networks, and we can connect a variety of local network technologies to it.

Our Cisco routers are capable of interconnecting a wide variety of media and of routing IP, IPX, Appletalk, DECnet, SNA, X.25, and most other protocols that concern us.

The academic network sees heavy use. During working hours, our routers are typically handling from 500 to 700 active IP connections.

3.2.4. Departmental Networks

Departmental computers are becoming increasingly networked. Macintoshes are on Appletalk networks, PCs are usually on Novell networks, UNIX workstations are on TCP/IP networks, and VAX/VMS systems are on DECnet and/or TCP/IP networks, or connected host-to-host via BITNET. Departmental IBM mainframes are connected by BITNET and/or TCP/IP.

The Health Sciences campus has large Token Ring and Ethenet networks, as well as Appletalk, DECnet, Arcnet and other types of LANs, many of them attached via SNA gateways to IBM mainframes. The Lamont campus has a large TCP/IP network comprised mostly of UNIX workstations. The Engineering

---

[4]Certain special cases don’t fit this model: leased-line BITNET connections (which are being phased out), dialup UUCP connections, etc. These are integrated with our TCP/IP network to various degrees through our central host computers.

[5]For Token Ring, Media Access Units (MAUs) are required on each floor, or every two floors.

School has a large TCP/IP network. Health Sciences, the Law School, and the Business School all have large Novell networks.

3.2.5. Access to Central Services

Access to the administrative network from the user's desktop is via:

1. Asynchronous terminal or emulator connected through the Rolm system, direct from an authorized campus Rolmphone, to an IBM 7171 for full-screen access to CUVMC. The 7171 performs 3270 emulation.
2. Asynchronous terminal or emulator connected through the Rolm system, direct from a campus Rolmphone, to a port on the COMTEN for linemode access to CUVMC (this access mode is being phased out).
3. True 3270 terminal or PC equipped with Irma (or similar) board and 3270 emulation software, connected by Rolm wiring or coaxial cable to a 3x74 control unit, thence to the COMTEN. A PC can furnish the 3270 screens to the user, or it can hide them via a screen parser such as Easel, or through the SNA 3270 API.
4. 3270 emulation via SNA from an IBM Token Ring network.
5. Program-to-program SNA protocols such as APPC or LU6.2 between PC LAN applications and central administrative applications.
6. SNA protocols, usually 3270-based, from VAX/VMS or other departmental computers.

File transfer is via Kermit, IND$FILE, or screen capture.

Desktop access to the academic network is via:

1. Asynchronous terminal or emulator connected through the Rolm system, either direct from a campus Rolmphone or by dialup, to an IBM 7171 for full-screen access to CUVMB. The 7171 performs 3270 emulation.
2. Asynchronous terminal or emulator connected through the Rolm system, either direct from a campus Rolmphone or by dialup, to a port on the COMTEN for linemode access to CUVMB.
3. Asynchronous terminal or emulator connected through the Rolm system, either direct from a campus Rolmphone or by dialup, to a terminal server, and from there via Telnet, Rlogin, or similar protocol to any of the academic host computers or services.
4. From a desktop workstation, direct access over the network via Telnet, Tn3270, Rlogin, LAT, CTERM, X, or other protocols.
5. For PCs and/or Macintoshes on a LAN, via a LAN operating system and shared file server, possibly gatewayed or routed to the backbone network.
6. Peer-to-peer IP-based protocols, including SUN Remote Procedure Call (RPC).

File transfer is via Kermit, FTP, or screen capture (or remote copy, or NFS, etc).

3.3. Expansion of the Academic Network

In September 1990, the Columbia University Network Architecture Task Force endorsed a plan for installation of a new fiber optic backbone campus network extending to all buildings on the Morningside campus and many of the nearby off-campus buildings. In August 1991, the funding plan was submitted to the University Provost and the ISS. The plan has been approved and the installation is in progress. This will be the Morningside campus common network of the 1990s.

The fiber project is being funded from AcIS resources. Detailed planning, design, and implementation is being done in AcIS at the Director level and by the various AcIS groups and with coordination and cooperation with other University groups (Communications Services, the Law School, Facilities Management). The backbone network will be owned and operated by AcIS.

The CLIO terminal network is being converted from dedicated point-to-point circuits to Ethernet, TCP/IP, and Telnet/Tn3270, using the backbone.

3.4. Security

Network security depends largely on a common network that is difficult or impossible to tap or spy upon. The architecture of Columbia's common network approaches this goal. The fiber optic backbone is virtually impossible to tap. Within buildings, sensitive administrative applications can go through a separate network hub.

Security also depends on host access procedures, human factors, and so on. With these items properly accounted for, the common network can provide secure access to administrative applications.

A joint University–Hospital committee on on computer and data security is studying these issues and will make recommendations to the University Provost and President of the Hospital. Security is discussed more fully in Appendix II.

4. Migrating to the Common Network

Consistent access to administrative and academic resources from each desktop computer or terminal is provided most effectively when a common network and set of protocols can be used to access all central resources.

The academic backbone network, especially after it is expanded to include all Morningside campus buildings, should be the common network. The preferred common transport/network protocol is TCP/IP.

A common network and protocol reduce our overall expenses and commitments, are more easily and effectively managed, and expand and simplify access for our users.

What is required to bring the administrative systems onto the common network? Academic applications fit the common network model. All the central UNIX and VM/CMS hosts and services can be accessed via Telnet or Tn3270 or higher level protocols built on TCP/IP. The remainder of this section examines the administrative applications.

4.1. Alternatives

Today's segregated administrative and academic networks pose fundamental problems for users and the University as a whole.

A significant class of users needs access to both administrative and academic data, and often needs to integrate the two types of data together. Currently, this is difficult or impossible. Ideally, an authorized user should be able to access both types of data simultaneously on the same workstation, for example in different windows: cutting budget or personnel data from a window on an administrative application and pasting it into an academic window containing a research proposal to a funding agency. Similarly, faculty members whose workstations are normally used for research should be able to learn the status of their research grants at the same workstation, often resulting in losses or overruns that could be prevented by convenient online access.

Keeping the networks separate also imposes a burden on the University as a whole in unnecessary expenses for duplicated resources: cabling, equipment, and people.

Our alternatives are:

1. Maintain the status quo: separate networks. This is undesirable because of the ongoing costs in money, human resources, and insufficient service delivered to the users.
2. Convert the academic network to SNA protocols. This is impractical because most academic computers and terminals do not (and can not) support SNA.
3. Convert the administrative network to ISO OSI protocols. This is not possible because the required software and services do not yet exist.
4. Convert the administrative network to TCP/IP protocols as cost, performance, and capacity allow. The software and technology are readily available.

The remainder of this section explores alternative number 4.

4.2. Present-Day Applications

Most, if not all, of today's administrative applications are based on the 3270 terminal model. These application could also be accessed on a TCP/IP network via TN3270 protocol if TCP/IP were available on the administrative mainframe. TN3270 programs are available for UNIX computers and workstations, PCs, and Macintoshes, and TCP/IP is available for VM and MVS. The primary online administrative applications fully accessible in 3270 screen mode are listed below.

ARGIS	Alumni Development System
OPG	Project and Grants Tracking System
LDS	Labor Distribution Inquiry System
Payroll	Payroll Inquiry System
BIS	Benefits Information System
FAS	Financial Accounting Inquiry System

CAPS	Automated Purchasing System
APS	Accounts Payable System
FAM	Financial Aid Management System
SOSIS	Old Student Information Inquiry System
SIS	New Student Information System, except for file transfer.

Several existing or forthcoming administrative applications require SNA-specific technology:

IRM	The Imaging Record Management system for financial aid requires SNA PU2.1 and LU6.2 over Token Ring.
SIS	New Student Information System, file transfer (PCI feature) requires 3270 hardware (such as an Irma board) in the PC.
FOCUS	The Distributed FOCUS (not currently in use) file transfer feature requires 3270 hardware in the PC.
CDR	Transission of Rolm CBX Call Detail Recording (CDR) billing records to WCS billing software from NetView/PC to CICS/DDM via LU6.2.
IBMLINK	Access to IBM's online Customer Support system, IBMLINK, requires an SNA PU Type 4 connection.
COBOL	CASE tools like MicroFocus Cobol require the IBM 3270 emulation API and IND$FILE for its integrated source code file transfer function.

Vendors of these applications have been or will be encouraged to support more conventional access methods. In any case, except for IRM and SIS, these are not end-user applications and their dependence on SNA does not pose a serious obstacle to the idea of a common network, because SNA can operate over the common network (see Appendix I).

4.3. Future Applications

New applications depending on the SNA 3270 API, APPC, LU6.2, IND$FILE, and other SNA or other proprietary protocols should be avoided. We have informed our major actual and potential vendors of administrative application software of our direction, and have received encouraging responses.

4.3.1. AMS

The Network Integration Task Force met with Mike Titmus and Robert T. Lindsay Jr of AMS on December 16, 1991. Robert's title is ''Senior Principal'' — he is a top architect of AMS software design, setting the direction for the entire company, not just one product line or marketplace.

We have contracted (or will contract, or are considering whether to contract) with AMS for our major administrative applications:

IRM	Imaging Record Management system for financial aid
SIS	The Student Information System (currently installed)
CUFS	The College and University Financial System

HRS The Human Resources System

The last three are AMS "core-based" systems, meaning that they are designed for portability using a layered architecture, with separations at transport, data, transaction, and presentation. Current AMS applications use an APPC transport, but can be ported "in a matter of weeks" to a TCP base as soon as a customer requests it. TCP transport is already in use in their government applications.

IRM is a joint AMS-IBM development communicating over an APPC transport, and is not necessarily designed for portability or easy conversion to a TCP base.

The AMS core applications operate via several different user interfaces:

1. 3270 screens via VSAM and CICS. This mode can be supported by the current 3270 network, via async terminals or PCs with Kermit thru 7171s or other 3270 protocol converters, or by Tn3270. SIS works in this mode; we have not contracted for a client-server version of SIS.
2. X presentation. Requires an X server on the user's workstation, but beyond that, no cooperative processing. Any platform with an X server can be used, even PCs.
3. Distributed server/client applications, with central or distributed databases over any mixture of SNA and TCP connections. Client software will be available for PCs with Windows, Macintoshes, UNIX and workstations (RS/6000, SUN, etc). They call this their "MicroTradeLine architecture."

Robert expressed AMS's formal commitment to operate their core applications over a TCP base within any time frame we request.

Robert was asked whether he knew of any reason why we should not adopt TCP/IP as our campus standard as far as AMS products were concerned. His reply was "no." Of course he could not speak for other vendors and predicted that SNA-specific applications (not from AMS) would be around for a long time.

Robert identified OSF DCE (Distributed Computing Environment) in general, and Kerberos authentication in particular, as a corporate direction, but did not make a firm commitment. He stated that end-to-end encryption would be done at the transport level. These measures should be sufficient to make administrative applications secure in a shared-network environment where tapping is possible, provided the applications are configured and located appropriately. He could not say whether or to what extent AMS would make use of DCE's access control list facility. AMS has no plans to incorporate "physical security" measures (e.g. fingerprints, voice recognition, ID cards, etc) into their products.

4.3.2. Easel

Easel, whose product is used to create "user-friendly" PC applications based upon 3270 connections via screen parsing and key mapping, has also been informed of our direction.

Easel software is dependent on a 3270 emulation API called HLLAPI (High Level Language Application Programming Interface). According to Andy Ellicot, a technical person at Easel, Easel does not plan to directly support TCP/IP or TN3270. However, they do support any emulation product that provides an

IBM-compatable HLLAPI. Andy believes that ICON Technologies has a Windows 3270 emulation product that supports TCP/IP and has a HLLAPI interface. In addition he has heard that Attachmate Corp is working on adding TCP/IP support to their Windows 3270 emulation products. He thought it might be available sometime in the first quarter.

Both of these approaches would greatly improve our ability to utilize API or IND$FILE dependent software over TCP/IP. This would include Easel, Micro Focus Cobol and probably any local code written to the 3270 API.

In the meantime, Easel has been modified at Columbia's behest to support asynchronous connections to 3270 protocol emulators such as IBM 7171s or Cisco terminal servers equipped with tn3270.

5. Network Design, Operation, and Organizational Structure

The Columbia University academic network has evolved from a small experimental Ethernet into a large, serious production network used by thousands of people. Network outages occur with some frequency and sometimes result in disruptions for those dependent on the 24-hour availability of the network and network-accessible applications.

And yet, there is only an informal network operations function. With the integration of administrative functions into the common network, the need for a reliable and responsive network becomes more critical than ever. The components of such a network function include:

- Coordinated design and planning
- Coordinated installation and support
- Coordinated configuration management
- Monitoring and management tools and the people to run them
- An organization structure that ensures the rapid (or at least predictable) isolation and correction of problems (faults, security violations, congestion)
- A funding mechanism that allows the network to grow and improve as needed

5.1. Design and Planning

The design of the network — its physical and logical configuration — must be a "known quantity". Weak points should be identified, and if they are critical, they should be strengthened: fixed, upgraded, or redundant components installed. Migration to new technology must be accomplished in a coordinated way. Design and planning functions include:

- Needs and financial analyses
- Research and development
- Standardization
- Upgrades, testing of new components

- Backup circuits, redundant components, dynamic reconfiguration
- Migration to or addition of new technology and/or protocols

5.2. Design and Installation of Departmental Networks

There should be one place — one office, one phone number — for users to order network installations that fit with our overall connectivity strategy and can be adequately supported. Since the backbone network is owned by AcIS, connections of local area networks to the backbone must be coordinated and approved by AcIS.

Other issues must be considered here. The initial interview, needs analysis, etc, must be done by people who are familiar with the user's applications, LAN technology, and the backbone network.

How or whether a combined AIS/AcIS organization provides LAN support to the University's departments and offices is an issue for further study by the LAN Support Task Force.

5.3. Configuration Management

There must be a central authority to decide which protocols can be supported (and at which levels) on the backbone network. That authority presently resides within AcIS, which provides University-wide network connectivity services on the academic backbone.

A viable network requires a central registration authority for the names, numbers, and addresses of network objects: host names, host addresses, mail domains and addresses, etc. TCP/IP, Appletalk, Novell, and DECnet names and addresses are managed by AcIS, with some authority being delegated to the subnetworks. SNA names and addresses would be managed by AIS on Morningside and CIS at the Health Sciences Campus.

The Morningside academic backbone network consists of a collection of routers connected by the backbone cabling, with connections to hubs in each building. The configuration of the routers and hubs is critical to the correct and secure operation of the network. The responsibility for router and hub configuration lies within AcIS.

The responsibility for configuration of other network components — host computers, front end processors, protocol conversion devices, terminal servers, various other types of servers, etc — lie within various groups that own or control these devices. Each of these devices can affect all the others, and the network as a whole. Thus all network configuration activity and information should be closely coordinated.

5.4. Monitoring Tools

UNIX-based or other network monitoring and management workstations can be placed in strategic locations to detect and isolate faults. A wide variety of network management software is available to perform these functions. Open network management protocols (SNMP and CMIP) should be preferred over closed or proprietary ones. Most current network management software is based on SNMP, TCP/IP's Simple Network Management Protocol and/or IBM (SNA) network management mechanisms.

It is desirable that our network monitoring tools extend as far as possible beyond the backbone network, preferably all the way to the user's desktop computer: its network interface, its network control software, even its application software.

Network management software should be chosen that provides the widest and deepest coverage of all aspects of our network. This, in turn, requires that as many components of the network as possible are ''manageable'', i.e. provide management information.

Our network management software should have the following capabilities:

- A selection of views of the network
- Detection and isolation of faults on the backbone (and beyond)
- The ability to turn off misbehaving network components and turn them back on after they are fixed, to include entire network segments, specific router interfaces, specific workstations, printers, modems, etc.
- Security management: prevention of, and/or the ability to detect, unauthorized network access
- Traffic analysis
- A database of network components, preferably distributed to make different groups and departments responsible for their own pieces, and possibly including technical specifications and documentation; contact, ordering, and configuration information, etc.
- Discovery of network objects not in the database
- Accounting management
- Remote configuration capability
- Trouble ticket management
- Report generation, including trend analysis
- Compliance to open standards (SNMP, CMIP, OSF/DME)
- Extensibility to incorporate network objects or functions not directly supported by the vendor

The network management system should fit the server/client model, so it can be used from various locations, rather than only on a dedicated PC. The presentation interface should be open, to permit its use on a variety of platforms, for example using X, Motif, or similar regime for client/server communication.

Today's top contenders are Cabletron Spectrum, which is already operational at Columbia and has been configured for the backbone network as well as many of our LANs, and IBM NetView/AIX. NetView,

unlike Spectrum, can handle the Rolm CBX and the SNA network, but cannot handle our Cabletron equipment (although NetView can accept SNMP traps).

5.5. The Network Operations Function

A formal network operations function will be established within AcIS for the academic network. AcIS staff will set up network monitoring and management stations and place them at various strategic locations. Personnel from AcIS — and hopefully other groups — will be appointed to monitor these stations (AcIS is in the process of establishing two staff positions for network operations). These people will constitute the new Network Operations Group. Among the responsibilities of this group:

- At the network management workstation, monitor faults, errors, traffic, etc.
- Handle trouble reports.
- When a fault occurs, gather as much information as possible about the fault, log the fault, and fix it or else contact the appropriate group(s) to have it fixed.
- Follow up on trouble calls: make sure the user is kept informed, inquire periodically about unresolved faults, etc, until the log entry is closed.
- Produce periodic reports on network usage and performance.

Coordination among all groups involved is critical.

5.6. Organizational Structure

Network planning, installations, configuration, upgrades, operations, monitoring, access control, and maintenance affect the University as a whole and require involvement of various groups from AIS, AcIS, and Communications Services, and to some degree the schools and departments. A management and reporting structure that crosses group boundaries is required.

For network operations, there must be a clear path for problem resolution:

- Who is responsible for fixing physical connectivity problems (broken physical connections, tripped circuit breakers, etc), noting that there can be different contact points depending on the location and nature of the problem: Morningside, Health Sciences, Barnard, Teachers College, Lamont, NASA, EMEC, South America, etc.
- Who is responsible for fixing configuration (e.g. routing, congestion) problems.
- The responsible person for each departmental network.
- Who to call for wide area network problems (the phone company, a local network service provider such as PSI, etc).
- Who to call when the nature of the problem can’t be determined.
- Access procedures: the network operations group should be responsible for ensuring and coordinating access to all the locations on the Morningside campus where our backbone equipment resides, including building basements, tunnels, etc.
- Hours of coverage should be clearly established for each area of responsibility.

- Emergency procedures should be established for critical events that occur outside the hours of coverage.
- Contacts for security events.
- Escalation procedures for all types of faults should be put in place.
- A mechanism for ensuring that all relevant groups do their parts and cooperate with each other is essential.

For network installations and upgrades, the organizational structure must also be clear:

- Who is responsible for design and installation of the backbone network — cabling, routers, and hubs. This is presently AcIS.
- Who is responsible for the configuration of the various components of the network: routers, front ends, central host computer systems & servers, etc. AcIS is responsible for routers and other equipment directly on the academic backbone.
- Who is responsible for connecting departmental LANs to the backbone network, including advising departments on the recommended LAN technologies.
- Who is responsible for helping departments with localized LAN problems.

This implies an organization structure that ensures that people from different groups come together periodically to review and refine the areas of responsibility:

- Who ‘‘owns’’ what
- Precise demarcation points
- Resolution of disputes

Finally, how do the users see us? Should there be one central help desk (or networks help desk) to receive calls from all users and route them appropriately, or should there be separate desks for different topics and/or different classes of users?

6. RECOMMENDATIONS

The administrative systems should move to the common network and convert from SNA to TCP/IP. This can be done in stages. The SNA and TCP/IP protocols can coexist for some period of time on the common network.

6.1. Physical Network

There should be a common physical network for administrative and academic use, and that network should be the AcIS backbone network combined (for buildings connected to the Rolm system) with the Rolm building wiring. Ethernet, Token Ring, and Apple Localtalk can be used on the Rolm wiring. The common network is based on Ethernet, but can be migrated to another protocol (such as FDDI) in the future, transparently to the local networks attached to it.

The AIS network can be migrated to the common network to eliminate high-cost, low-performance dedicated point-to-point circuits and to achieve economies of scale, consolidation of effort, consistency, manageability, and other cost savings. New installations should use the backbone. Old installations can be converted over time for economic or management reasons. See Attachment I.

6.2. Network Protocols

Central, common resources on all our mainframe hosts and servers should be accessible via TELNET, TN3270, FTP, and/or other higher-level protocols built upon TCP/IP.

IPX and Appletalk should be supported on the common network to access any central or common resources that require these protocols, such as Novell file servers, Apple print or file servers, etc.

DECnet, SNA, and other protocols that can peacefully coexist with our supported protocols should be supported on the backbone as ‘‘guests,’’ meaning that efforts to integrate these protocols into our central network management systems and other central applications will be given low priority.

SNA and other closed, proprietary protocols currently in use should be phased out in favor of open protocols, and should be avoided in new applications.

A consistent, standard set of *open and interoperable* higher-level and management protocols should be adopted. For the present, these require a TCP/IP (or asynchronous terminal) base; in the future they can be migrated to OSI. Many of these are still on the drawing boards, so this list must be refined over time:

- Network management: SNMP (with migration possibilities to OSI CMIP or OSF/DME).
- Authorization and security management: To be determined; most likely OSF/DCE and/or GSS (Generic Security Service)
- File transfer: FTP and Kermit, plus OSI FTAM in the future.
- Virtual terminal service: Telnet, TN3270, rlogin (with migration possible to OSI VT).
- E-Mail: SMTP (with possible migration to or coexistence with CCITT X.400)

- Directory service: CCITT X.500
- Remote procedure call: To be determined, most likely OSF/DCE
- Distributed file system: To be determined (OSF/DCE, NFS, Novell, etc)
- Database access: To be determined by the data management task force (candidates include X/Open SQL/Access, SQL2, etc)
- Screen presentation: To be determined by the Presentation Task Force (candidates include VT-300, X, OSF/Motif, Open Look, Macintosh, Microsoft Windows, OS/2 Presentation Manager, etc)
- Character sets: ISO 8859-1 Latin Alphabet 1 (migrating to ISO 10646)

TCP/IP should be installed under VM on the administrative mainframe for an evaluation period as soon as a plan can be formulated. The primary considerations are:

- Performance of TCP/IP connections.
- Maximum number of simultaneous TCP/IP connections before system performance degrades unacceptably.
- Does each application work when accessed via TCP/IP?
- Security.

If the results are unsatisfactory, a plan should be made for resolving the problems that were identified. This would include working with IBM software engineers, perhaps under nondisclosure. We know that TCP/IP works well under VM/CMS from experience with CUVMB, so the major question is capacity. In the worst case, we can run with a mixture of TCP/IP and 7171 connections while waiting for IBM to improve the efficiency of their TCP/IP products.

TCP/IP, including DCE (beta), can be installed under MVS too. The major benefit would be an open method of file transfer (FTP) for our TCP/IP network users, in addition to Kermit. AIS must weigh the costs and benefits of MVS TCP/IP and decide for itself.

6.3. User Access

The following methods should be supported:

- Asynchronous terminal or emulator connected through the Rolm system (or equivalant data switch), either direct from a campus Rolmphone or (if it can be done securely) by dialup, to an IBM 7171 or other 3270 protocol converter (e.g. in the Cisco terminal server, or the Session Manager) for full-screen access to CLIO and the IBM mainframe.
- Asynchronous terminal or emulator connected through the Rolm system, either direct from a campus Rolmphone or by dialup, to a terminal server, and from there via Telnet, Rlogin, or similar network protocol to any of the academic host computers or services.
- From a desktop workstation, direct access over the network via Telnet, Tn3270, Rlogin, X, or higher-level TCP/IP-based protocols.
- For PCs and/or Macintoshes on a LAN, via a LAN operating system and shared file server, possibly gatewayed or routed to the backbone network.

- Supported connection technologies for departmental networks include Ethernet, Token Ring, and Localtalk. Supported network/transport protocols include TCP/IP, IPX, Appletalk, and (to be phased out) SNA. Departmental networks are always connected to the backbone through a router.
- File transfer via Kermit, FTP, or network remote file access.

6.4. Work Items for Other Groups

The following sections list issues raised in the Network Integration Task Force that are left to other groups better constituted to handle them.

6.4.1. Directors

The Directors of AIS, AcIS, and Communications Services, possibly together with higher-level officials, must decide the questions of organizational structure, reporting relationships, and finances, particularly with regard to network planning, operations, and management:

- What body makes decisions about future directions of the network — architecture, technology, protocols, standards? What groups, schools, and departments should be represented (such as AIS, AcIS, Communications Services, Health Sciences, Lamont, the Libraries, the Engineering School, CTR, Arts and Sciences, etc)?
- What is the organizational structure for network operations? The lines of authority, the points of demarcation, etc? Should a network operations committee composed of representatives of the different groups (and campuses) be constituted with the authority to make policy and resolve questions of an operational nature? Or should such a committee serve in a purely advisory capacity?
- Can we present a unified face to our users? A single point of contact, an effective escalation procedure?
- How is the network financed? What (if any) are the chargeback mechanisms, policies, and rates? How do funds flow among AIS, AcIS, and Communications Services? In what ways are users charged for connection to and use of the common network? Are administrative and academic users charged the same way?
- Given the move to a common network, who will control the existing SNA network?
- Who decides the scheduling of the fiber backbone installation, in particular the order in which buildings are connected? If the current plan meet does not the needs of AIS, what can be done to change it?

6.4.2. LAN Management Task Force

The question of how and which types of departmental LANs are to be installed, supported, serviced, is left to the LAN Management Task Force. Among the items to be resolved:

- Who deals with the departments, and on what basis (advisory, fee-based, etc)? Is there a single group or a multiplicity of groups? If more than one group, how are they to be coordinated?
- What services should be offered to departments: hardware resale, installation, configuration, management, backups, site licensing of software, software installation, software consulting, linking with central servers, user ID entry and management, etc?

- Which hardware platforms, networking technologies, operating systems, network operating systems, network protocols, applications programs, vendors of all the preceding, etc, will we support or condone, and to what degree?
- What types and levels of support can be offered?

6.4.3. The Data Management Task Force

The protocol to be used for accessing remote or distributed databases: SQL/Access, SQL2, etc.

6.4.4. The Presentation Task Force

Asynchronous terminal access to central resources should be supported indefinitely, and that implies supporting a specific repertoire of terminal types. We recommend, at minimum, the DEC VT-320 (international version) be supported.

Distributed applications for the user's workstation use the presentation functions of the workstation environment. These might include X, OSF/Motif, Open Look, the Macintosh environment, Microsoft Windows, OS/2 Presentation Manager, etc. The Presentation Task Force should make a list of approved or recommended presentation environments.

All presentation methods should support ISO 8859-1 Latin Alphabet 1, for the entry and display of the accented and other special characters of the Western European languages (English, French, Spanish, Italian, Dutch, German, Norwegian, Swedish, Danish, Finnish, Icelandic, Portuguese). In the future, we should expect to support the full range of the world's writing systems by migrating our equipment and applications to the forthcoming Universal Character Set, ISO 10646.

Appendix I. SNA to TCP Migration Options

Prepared for the AIS Network Integration Task Force by Alan Crosswell, AcIS, November 24, 1991, with amendments by Stew Feuerstein, AIS.

This paper is divided into two sections. The first briefly describes the current AIS SNA network and then outlines a number of approaches that are available in order to continue to support SNA in the context of a common university network.

The second section outlines some of the options and issues in replacing the functionality of SNA mainframe host and workstation (and control unit) software with TCP/IP.

Neither section presumes to answer all the questions, but I hope the issues raised can become a useful basis for further discussion within the task force.

I.1. Supporting SNA

First, an overview of the current AIS network, followed by descriptions of a few ways to continue to support SNA. Some of these are currently being done while others are available options.

I.1.1. Current State

The current AIS SNA network is quite new, therefore, not yet heavily entrenched in SNA-specifics such as LU6.2. There are two major ways that user terminals and workstations attach to the AIS VM and MVS operating systems — via IBM 7171's and the COMTEN 3690 Front End Processor.

AIS has two IBM 7171's, each capable of supporting 64 simultaneous RS-232 asynchronous ASCII terminal users at speeds up to 19,200 bits per second. These terminal ports are available on the ROLM data switch and are connected to by users on ASCII terminals — primarily PC's emulating a VT100 with Kermit. The 7171 performs 3270 terminal emulation so that the ASCII terminal looks like a 3278 display station. Additionally, the 7171 supports "transparent mode" which allows transfer of arbitrary ASCII data. Transparent mode is used for Kermit file transfers and can also be used by applications such as SAS/Graph to transmit Tektronics or pen plotter graphics commands.

A COMTEN 3690 front end processor (FEP), shared by AIS and AcIS, supports a vanishingly small number of ASCII "start/stop" or "line mode" terminal connections via the ROLM switch as well as synchronous communications for remote BITNET sites (on the AcIS side) and approximately fifty 3x74 cluster controllers along with a small number of RJE stations. With the exception of two recently-added 56 kilobit per second interfaces, all the cluster controllers operate at 9600 bits per second using SDLC (having recently converted from bisync).

The COMTEN is several years out of date and sorely in need of retirement or replacement as the hardware and software are rapidly becoming extinct. NCR/COMTEN has already informed us that support for the operating software will shortly move to "category II" — meaning any bugs will be fixed on a time and materials basis rather than being covered by service contract. AcIS has set a goal of being entirely off the COMTEN by the end of this fiscal year. This includes termination of start/stop service

and re-homing of bisync BITNET links onto TCP/IP (with Princeton's VMNET software) or by moving the connecting sites to other institutions.

Each 3x74 cluster controller supports somewhere between 8 and 32 coax-attached 3270 terminals, PCs with IRMA boards, or printers (including laser printers attached by Malibu coax converters). The exact inventory of these devices is needed in order to determine costs for any continuance or conversion to different technologies.

The two recently-added 56K interfaces to the COMTEN each connect a remote 3174 Token-Ring control unit. One of these 3174s is located on the UDAR Token-Ring LAN in the Interchurch Center and the other in the computer center machine room, to support Health Sciences and Watson/Philosophy Token-Rings that are bridged by CUnet. PC's running SNA software on these LANs communicate with the mainframe, primarily by doing 3270 terminal emulation.

I.1.2. Token-Ring SNA

There are three areas of concern in Token-Ring SNA:

1. Support of end-user terminals/workstations.
2. Mainframe connection to the ring(s).
3. Common network infrastructure to support interconnected rings.

I.1.2.1. End User Terminals/Workstations

UDAR and some Health Sciences campus users are now using Token-Ring SNA to do 3270 terminal emulation on the mainframe from their PCs. SNA peacefully coexists on the Token-Ring LAN with other LAN and WAN protocols such as Novell's IPX and TCP/IP. PCs are sometimes able to simultaneously run two or three of these stacks at the same time, subject to memory limitations and how well-behaved the software is.

Because these PCs are on a LAN rather than being attached to a control unit (e.g. with an IRMA), they are potentially able to do a lot of non-SNA mainframe things such as use LAN file servers, and run network applications.

Current 3x74 controllers and their terminals can also connect to a Token-Ring. Depending on the model of 3x74, this either requires adding an interface card or replacement of the controller. Converting a 9600 baud SDLC connection to Token-Ring will lead to substantial throughput improvements.[6] This could allow coalescing two 32-port 3174s to a recently announced upgrade which supports 64 ports.

Any PCs that were connected to the 3x74 have the option of remaining there or having a Token-Ring card installed (to replace the IRMA). "Real" 3270 terminals can only connect via a control unit.

---

[6] an IBM specialist has stated that the 3174 is constrained only by the speed of the connection to the mainframe; it has sufficient CPU capacity.

It would seem that the goal would be to have any PC directly on the LAN rather than connecting it to a control unit since the functionality of a LAN-attached PC is a superset. Continued support of coax 3270 terminals requires control units. The cost/benefit of replacing 3270 terminals with PCs needs to be assessed.

Finally, note that requiring end-user devices to support SNA Token-Ring eliminates a large percentage of potential workstations, such as those that support Ethernet. See the section of SNA Gateways, below for some options here.

I.1.2.2. Mainframe SNA Token-Ring Connection

There are a currently four ways to do SNA Token-Ring mainframe connections:

1. FEP SDLC to ‘‘Remote’’ Token-Ring 3174
2. Local Token-Ring 3174
3. FEP Token-Ring Interface
4. IBM 3172 SNA connections

As mentioned above, two remote SDLC 3174s of these have been installed, each at 56 kilobit/s. The 3174s that provide this function are fairly inexpensive (around $4000) and the COMTEN interface, while costly, was much cheaper than a Token-Ring interface for the COMTEN. One of the two installed 3174s is actually remote, located at Interchurch. The other is located in the machine room on a very short cable to the Comten (actually a CSU/DSU elminator — or null modem). Since the 3174 connects to the FEP, most SNA processing is still done in the FEP rather than in the mainframe host.

A different model 3174 can be channel-attached (‘‘local’’) to the mainframe (the current local 3174 cannot be used because it is configured as a non-SNA controller, which is required because it has operating system consoles on it), configured as an SDLC SNA controller, and equipped with a Token Ring interface board, at a cost of about $14,000.

A common complaint heard in discussions of using local 3174s for Token-Ring is that most of the SNA processing is now done in the mainframe since there is no FEP to offload to. However, this can be minimized by only configuring a small number of PUs on the ring.

Useful numbers to compare the performance impact of local 3174s to FEPs are sorely needed.

FEP Token-Ring interfaces are also available. This adds a 4 or 16 megabit/s Token-Ring interface to the FEP. Depending on the FEP, this can provide a very high-throughput connection, supporting a large number of end-user devices. FEP vendors claim that this method provides substantial offloading of CPU cycles from the mainframe when compared to a non-FEP solution such as a local 3174. However, especially an FEP and also the FEP token-ring interfaces require a fairly large investment.

In the case of the existing COMTEN, adding a direct token-ring interface was deferred due to cost and the fact that a remote SDLC 3174 would provide sufficent performance in the short term and could be easily recylced for other purposes. Also, the current COMTEN is severely performance-constrained in terms of

supporting added hardware. And, any Token Ring interfaces added to COMTEN would have to be replaced when the COMTEN is replaced.

I.1.2.3. Common Network Infrastructure Support for SNA

As already demonstrated for the Watson/Philosophy SIS and Health Sciences rings, the common network is able to provide Token-Ring source-route bridging, in this case using an arbitrary network cloud to provide the connections between the rings.

An additional level of SNA network support — for SDLC — is described in the following section. In addition, we are currently researching the capacity of Cisco SNA source route bridging for high-traffic applications like IRM.

I.1.2.4. SDLC and SDLC Tunneling

The largest cost component in COMTEN replacements that were presented recently in formal proposals from IBM and COMTEN had to do with the number of SDLC links needed (50) to support existing control units. Substantially smaller and less expensive FEPs are available from both vendors that provide Token-Ring (and Ethernet) connectivity if the number of slow SDLC interfaces can be reduced.

An interesting option that has come out of Cisco Systems recently (as of software release 8.3) is the "Serial Tunnel." Wellfleet also has done this. See the March 1991 issue of Data Communications. Following is an excerpt of mail sent to Doug Ingling in late April, 1991 when Cisco was looking for beta-testers for the product:

Cisco Systems has asked us to be a beta-tester for their "serial tunnel." The way this works is you connect one Cisco router (the same exact ones that we already use for the campus network) via SDLC to a FEP like the COMTEN or 3745. You tell the FEP that this line is multidrop. Then, you connect your current existing 3174s (once converted from BSC to SDLC) to serial ports on the same or other Ciscos on the campus network. These 3174 SDLC connections (typically running at speeds of 9600 or 19200) get routed over the campus network to the Cisco that is connected to the FEP where they all get multiplexed onto the *one* SDLC link to the FEP. The major implications of this are:

1. A much smaller FEP can be used since you have fewer lines going into it.
2. A common campus backbone network can be used to provide communications even for the existing installed base of 3174s.

I.1.2.5. SNA Gateways

There are a large and diverse number of SNA gateways. These are generally devices that convert some other protocol into SNA. For example, Novell makes such as gateway so that PC clients use IPX to talk to the gateway which then has an SDLC link to the mainframe FEP. Note that this works with any of Novell's supported network hardware (i.e. not just Token-Ring). Similar solutions are available, for example in McGill University's NET3270, Digital Equipment's DECNET/SNA gateway, Sun Microsystem's Sunlink/SNA, etc.

One may also roll their own as Soumitra Sengupta ("Sen") of Medical Informatics and Computer Science has done by building a TCP/IP remote procedure call interface on top of SNA Services for the IBM RS/6000. Sen has written TCP/IP applications which do LU6.2 CICS transactions.

We have access to installed, working versions of the DEC and IBM products if anyone is interested in a demo (these are connected to the hospital SNA network).

I.2. TCP/IP as an Alternative to SNA

The following section describes some uses AcIS has made of TCP/IP with our VM and MVS systems and a few pointers to other things I've heard of. We do not have a good idea of how well the mainframe TCP/IP implementation scales to large numbers of users or for how its resource consumption compares to SNA. We need to check with other sites to get a better idea.

I.2.1. Mainframe TCP/IP Connection

AcIS has used an increasingly more powerful series of hardware to implement our TCP/IP connectivity. This started with IBM's DACU (Device Attachment Control Unit), which was a DEC UNIBUS controlled by a PC! We used an Interlan Ethernet controller. The DACU was followed by the IBM 8232 Lan Channel Station — a channel-attached PC/AT with standard PC LAN cards such as the Ungermann-Bass NIC. IBM followed the 8232 with the 3172, a channel-attached PS/2. Again, the 3172 also supports standard LAN cards. We chose Ethernet, but Token-Ring is also available. We currently own four Ethernet interfaces: a two-port 8232 (two PC/ATs in one box), and two Bus-Tech Inc. ELC-2 Ethernet controllers. BTI has recently announced a Token-Ring version of their product as well. The BTI's outperform the IBM 8232 and 3172 and cost less. The 8232 is available for loan or purchase by AIS and would certainly be an excellent platform for experimenting with TCP/IP.

While there are two or three mainframe TCP/IP software products available, we have only used IBM's TCP/IP for VM. We have also researched TCP/IP for MVS but decided not to purchase it at the time due to budget constraints and our ability to achieve the desired functionality with VM TCP/IP (since MVS runs under VM).

COMTEN also offers an Ethernet TCP/IP product along with their own FEP and host software. With this product, the FEP acts as the virtual terminal (TELNET) server, presenting the 3270 screens to the host in the exact same way as SNA terminals. IBM has also introduced an Ethernet interface for the 3745, but it is simply another way of doing a 3172 and does not at this time offload any processing into the FEP. It is not yet clear whether this product supports any form of file transfer (FTP or IND$FILE), which is a requirement for the administrative network.

I.2.2. Terminal Emulation

The TCP/IP virtual terminal protocol, TELNET, is an extensible protocol that has had various features added over the years. One of those is 3270 terminal emulation, frequently called tn3270.

For Unix workstations, we use the tn3270 program that is now standard with many Unix implementations. tn3270 emulates any of a number of 3278 models based on the workstation terminal emulation window size (e.g. 44 rows by 80 columns).

For PCs, we have used IBM TCP/IP for the PC (a "feature" of TCP/IP for VM) and Clarkson TCP (from

Clarkson University and based on NCSA Telnet from the National Center for Supercomputing Applications of the University of Illionois at Urbana-Champaign). We currently use Clarkson TCP, primarily because it is free, includes an FTP feature, and works with packet drivers.

On Macs we use NCSA Telnet for the Mac, with tn3270 support, including 3279GX emulation, added by Brown University.

For plain ASCII terminals, we are currently working on a Unix-based Session Manager, which will include tn3270 sessions. We hope to use this as the basis for a 7171 replacement for the VM and library systems — eliminating the ‘‘single host’’ connectivity of the 7171s. Cisco has also announced tn3270 support in their 8.3 software release for terminal servers.

I.2.3. File Transfer

Peer-to-peer file transfer between PCs and the mainframe is accomplished with the File Transfer Protocol, FTP. Clarkson TCP includes an FTP server on the PC end, so that the host can initiate transfers to and from the PC and vice-versa. The same holds for Unix and Mac TCP/IP implementations.

No direct support of IND$FILE file transfer is available in the free TCP/IP’s that we use, but it might be available in one of the commercial products — if it’s really deemed desirable to stick with proprietary file transfer mechanisms.

I.2.4. Printing

3270 users perform two types of printing: local (or control unit) printing and mainframe host-initiated.

I.2.4.1. Local Printing

The analogy to local control unit print for PC’s is hitting the Screen Prt key captured with Novelle print redirection. (There is no similarity to any print commands that go to PRN:).

I.2.4.2. Host Printing

For host printing, we use LPD, the de-facto standard TCP/IP print queueing protocol, based on Berkeley Unix’s Line Printer Daemon, LPD. While LPD is not terrific (only one printer per queue, no operator features like backspace, etc.) it does work in a distributed network environment. We have LPD client and server implementations under VM which are accessible to MVS as well by standard MVS to VM JES/RSCS spooling. The LPD code was written by Vace Kundakci under release 1 of VM TCP/IP. Release 2 includes a different LPD implementation now supplied and supported by IBM. A followon standard to LPD, MIT’s Palladium, was recently chosen by the OSF so we may see this emerging as the next TCP/IP printing standard.

I.2.4.3. Connecting Distributed Printers

RS-232 serial and parallel printers can be connected with TCP/IP terminal servers. A recently-announced gadget ($600 list price) is designed specifically for printers and will connect two serial and one parallel printer to an Ethernet. In both of these cases, a host computer is still required to control the printers. We use Unix hosts for this with our mainframe VM and MVS systems as clients. As mentioned above, we

run an LPD server on VM for IBM mainframe specific printers (e.g. machine room line printers, Xerox 9700 spooler).

Some printers are available with Ethernet and other network interfaces, essentially eliminating the terminal server from the equation and generally speeding up the printer due to the higher bandwidth connection. Many current Ethernet printers do not support TCP/IP but rather Novell printing.

Since most people on a LAN want both LAN and mainframe printing to come out on the same LAN printers, some way of making the mainframe a client of LAN printing protocols (such as Novell's) must be devised. Although we haven't tested it yet, I believe we have such an option available with a public domain PC-based LPD server — which should also be able to act as a Novell printing client).

The area of integrating LAN and mainframe printing — for TCP/IP and SNA — needs to be further explored.

I.2.5. Distributed Computing

LU 6.2 APPC is the basis of the SNA world's distributed computing architecture. A subset of the functionality as LU 6.2 is available as part of TCP/IP and there is a rich Distributed Computing Environment (OSF DCE) being built on top of TCP/IP or using TCP/IP as the stepping stone to OSI networking. DCE will contain equivalent and additional functions beyond those of APPC. Most TCP/IP distributed computing appears to be Unix based while LU 6.2 applications are DOS and OS/2 based. This is still a very new area and one that we are not particularly familiar with; Most of our use of computing has been terminal-to-mainframe based.

The IBM TCP/IP products are considered among the best TCP implementations when it comes to completeness of their implementations. For example, VM TCP/IP includes support for:

TELNET	virtual terminal.
FTP	file transfer.
SMTP	electronic mail.
NFS	network file system.
RPC	remote procedure call libraries.
X	X windows, including an X-windows client for GDDM.
Kerberos	network user authentication.

All or most of the above are also available with the MVS product.

I.2.6. Performance

There are several performance/capacity issues that will need to be explored in the migration from SNA to TCP/IP. Many shops are running TCP/IP on IBM mainframes, however we have not yet been able to identify any running five hundred simultaneous users. There are some concerns that the current IBM implementations of TCP/IP would perform poorly and consume significant resources, if we were to have all of our users connected via TCP/IP. We feel that IBM and other vendors will come out with new implementations that will greatly improve performance and reduce the resources used. Until that time we recommend a phased implementation with a period of SNA and TCP/IP coexistence. The duration and extent of the phased approach would be determined by cost/performance analysis and the anouncement of better performing TCP/IP implementations.

Appendix II. Security

Prepared for the AIS Network Integration Task Force by Alan Crosswell, AcIS, November 24, 1991

Network Security is an extension to traditional mainframe computer data security with the added complexity of physically distributed computers and data, and usually with the further concern of distributed ownership and administration of these systems. Following is a brief overview of security threats and some techniques used to combat them. It is important to note with network security as with any data security — even if the data storage medium is paper — that there is no black and white; a risk analysis must be performed based on such factors as the cost of implementing a given level of security versus the liability for disclosure or modification of data.

II.1. Definition of Terms

Spying

Spying is done by a computer that is connected to a network such as Ethernet or Token Ring. In the case of Ethernet, the computer puts the Ethernet adapter into so-called promsicuous mode at which point any and all data packets transmitted across the Ethernet are visible. Token Ring has an equivalent feature and is therefore no more immune from this kind of attack. Standard off-the-shelf PCs as well as network diagnostic equipment can be used for this. Another term for this is tapping.

Spoofing

An attacker spoofs a computer on the network by claiming to be another. For example, a spoofer might claim itself to be the network file server, mainframe host, or a particular client workstation when it is in fact not that computer.

Hacking or *Cracking*

This refers to the process of breaking in to a computer which can be accomplished in a number of ways including guessing legitimate users' passwords, exploiting operating system software bugs, etc.

Encryption

Encryption employs an algorithm to transform "cleartext" into "ciphertext" — that is, something that can't be understood. It is a useful tool for storing and transmitting data that might fall into the wrong hands. Various encryption techniques exist, the most well-know being the US Data Encryption Standard (DES). With DES, a secret key (password) is provided to the encryption algorithm in order to transform the data in ciphertext and the same key can be used to decrypt the data back to cleartext. The data is as secure as the secret key — which implies that the secret key should not be stored where it might be stolen.

Principal

Denotes a person or software entity (e.g. a server process).

Authentication

The process by which a principal identifies itself. For example, you authenticate yourself by typing your user ID and password. This proof of who you are is referred to as your credentials.

Authorization

The decision of whether a principal has permission to gain access to a resource (such as a host computer, application, datum). Upon presentation of your credentials, the application you are using decides whether you are allowed to perform the action you have requested. For example, having logged in you know issue a command to look up a record in a database.

Secret Key

A secret key or password is simply a secret shared among those "principals" (people, application processes) that need it. Your logon password is a secret key that is shared between you and the host operating software.

Public Key

Public Key encryption uses a unique combination of public and secret information, called your public and private keys. Your public key is something that can be freely disseminated, while your private key must be kept secret. If somebody wants to encrypt some data for you to later decrypt, they encrypt it using your public key. Once encrypted, the data can only be decrypted with your private key. This allows for principals to securely share data without having to maintain shared secrets. Everybody knows (can find out when needed) each other's public key and only the principal knows his private key, solving the problem of how to securely communicate a shared secret key by eliminating the need to.

II.2. Intra-LAN Threats

In the common University network, LANs are separated from each other and the backbone by routers. These LANs are generally restricted to a single building or department.

From within the LAN (that is, physical access to the LAN had been obtained):

- Spying may be performed trivially.
- Spoofing is also fairly easy, depending on the exact nature of the attack.

Threats from outside the LAN include:

- Hacking into a computer on the LAN and thereby gaining physical access. This can be done for example with multi-user workstations that allow network login such as telnet but can even be done where DOS-based products such as Carbon Copy are employed. These kinds of attacks generally use password cracking techniques to steal a legitimate user's password.
- Hacking a network server protocol such as NFS or Netware. These attacks are possible largely due to the original protocol designers' ignoring the possibility of hostile users on the network and/or laziness.

II.3. Inter-LAN Threats

The campus backbone provides Inter-LAN connectivity only. As such it has not computers connected directly to it other than the network routers which interconnect LAN subnets. Spying and spoofing can be done where physical access is available. Detection of attachment of a device to the backbone is subject to the technology used — copper wire is fairly easy to non-intrusively tap, while fiber would have to be broken and have electronics inserted into it. (The NSA is reputed to have a device that collects the light leaking out of a bent fiber as a means of non-intrusively tapping.)

II.4. Wide Area Threats

In a wide-area network environment it is best to assume that any and all data traversing it is publicly visible. While this may not be the case, no public network provider can make any guarantees about the secrecy of data traversing its links.

II.5. Techniques to Combat these Threats

While the above-mentioned threats are real, a risk analysis will indicate that most of them are trivial given the type of data traversing the network and other non-computer methods of disclosure and modification of information (e.g. employee abuse of the trust placed in them).

II.5.1. Password Protection

The largest threat by far is that of password cracking since an attacker may then logon and assume the identity of a legitimate user. If the cracker gets into a system administrator's account he may then be able to turn off audit trails, etc. For an entertaining account of how this is done, read Clifford Stoll's *The Cuckoo's Egg*, Doubleday, New York (1989).

A survey done at Bell Labs in the 1970's found that 30% of user passwords on one of their internal systems could be guessed trivially. So the first step is user education in selection and protection of passwords coupled with password management software that at least protects against picking truly easy-to-guess passwords. The Top Secret software on the AIS mainframe is an example.

Current password systems have one flaw when being used in a network login environment: even well-chosen passwords are usually transmitted in cleartext across the network[7]. For example, if you TELNET or start an SNA 3270 session to a host and then respond to the login and password prompts, someone who is able to spy on your session is able to capture your password as it goes by.

A software solution to this can be found in MIT Project Athena's Kerberos authentication system. Kerberos uses encryption to prevent cleartext passwords from ever traversing the network. Of course, if a password has been guessed, Kerberos doesn't help. Both client and server principals are authenticated with Kerberos, thereby preventing spoofing as well.

Other (more costly) solutions include phsyical devices such as smart cards. These are cards that a principal has in hand and are used as part of the authentication process. For example, the SecurID card from Security Dynamics has an LCD display with a unique random number that changes every 60 seconds. In addition to a conventional password, a SecurID user has to enter the current number displayed on the card, which is then checked with the SecurID server. If a card is lost or stolen, it can be invalidated on the central server, rendering it useless.

---

[7]One exception is the Novell server access password, which is transmitted from the client station to the server in encrypted form

II.5.2. Data Encryption

Data Encryption can be used to prevent disclosure of sensitive data. There are a number of ways encryption can be used:

- Files can be encrypted before transmission across the network and/or stored encrypted on a file system that may not be considered secure against attacks.
- Session encryption can be automatically done by network software. For example, a "Kerberized" Telnet implementation for Unix, besides supporting Kerberos authentication also allows one to turn session encryption on and off at will.
- Data link encryption can be done in data communications hardware, however, this is quite expensive and may not perform well enough at Ethernet speeds.

II.5.3. Auditing

No data security system is complete without audit trails. While the password and encryption techniques above can substantially reduce the threat of a break-in, it is still the case that the weakest link in the network security is the human being at the keyboard. Auditing is a requirement both for detection of and recovery from break-ins but also to protect against legitimate users simply making mistakes.

II.5.4. Commercialization of Open Network Security

Many of the security techniques discussed above are or will soon be commercially available. These include:

- Kerberos support at various levels is available from many vendors including IBM's TCP/IP for VM.
- The Advanced File System (AFS) from Transarc corporation is a Unix network file system product that is "kerberized."
- Digital and Sun have kerberized the Network File System (NFS).
- The Open Software Foundation's Distributed Computing Environment (DCE) includes Kerberos authentication and encryption of remote procedure calls as well as the Advanced File System.
- The Internet Engineering Task Force Common Authentication Technology working group is specifying standard authentication and encryption options for the TELNET protocol. These options include both secret-key systems such as Kerberos and public-key systems. This will lead to commercial implementations.

II.6. Policy Committee

The University and Hospital have a joint University-Hospital computer and data security committee, chaired by Steven Shea of the Center for Medical Informatics. This is a high-level policy-recommending committee convened by the Provost of the University and President of the Hospital. It is the appropriate forum to address network security policy questions.

II.7. Network Security Research Project

With support from AcIS, CIS, and Digital Equipment Corporation, a three-year research, development, and implementation project for a network user authentication and authorization system began on July 1, 1991.

The project is experimenting with and developing open network security techniques which can then be quickly turned over to production use in the common network.

Appendix III. Acronyms and Buzzwords

3174 An IBM control unit

3270 An IBM terminal

3274 An IBM control unit

3278 An IBM terminal

7171 An IBM 3270 protocol converter

AFS Advanced File System, a shared file system similar to NFS

AIX IBM's version of UNIX

API Application Programming Interface

APPC IBM's Advanced Program to Program Communication protocol

Appletalk Apple's datalink protocol

Arcnet A proprietary network from Datapoint Corp.

BDF Building Distribution Frame (Rolm)

BITNET An RSCS-based network for IBM mainframes and other computers

BSC Binary Synchronous datalink protocol

bps Bits Per Second

CBX IBM/Rolm's Computerized Branch Exchange

CCITT The International Telegraph and Telephone Consultative Committee of the International Telecommunication Union

CICS IBM's Customer Information Control System

CLIO Columbia University Libraries Information Online

CMIP The ISO OSI network management protocol

CMS IBM's Conversational Monitor System for VM

COMTEN A front-end communications processor used by Columbia's IBM mainframe

CSU/DSU Channel Service Unit / Data Service Unit for T1 circuits

CTERM The DECnet virtual terminal for wide area network connections

DACU Device Attachment Control Unit, an early Ethernet interface for IBM mainframes

DCE OSF's Distributed Computing Environment

DCM Data Communication Module for Rolmphones

DECnet DEC's proprietary networking method

DES The US Data Encryption Standard

DME OSF's Distributed Management Environement

DOS Disk Operating System

Ethernet A local area network technology in which stations communicate with each other at 10 Mbps over a shared cable in bus topology

Ethertalk Appletalk protocol for Ethernet

FDDI Fiber Distributed Data Interface

FDF Floor Distribution Frame (Rolm)

FEP Front End Processor

FTAM ISO's File Transfer and Management protocol

FTP The TCP/IP File Transfer Protocol

GSS Generic Security Service

GUI Graphical User Interface

HLLAPI IBM's High Level Language Application Programming Interface

IETF The Internet Engineering Task Force

IND$FILE An IBM file transfer protocol for use in the 3270 terminal environment

Internet The worldwide TCP/IP network

IP Internet Protocol, the network layer of TCP/IP

IPX Novell's Internetwork Packet Exchange protocol

Irma DCA's 3270 emulation product for PCs

ISO The International Organization for Standardization

Kbps Thousands of bits per second

Kerberos A TCP/IP-based security service

LAN Local Area Network

LAT DEC's virtual terminal protocol for local area (Ethernet) networks

Localtalk Appletalk protocol for twisted pair wiring

LU An SNA Logical Unit

MAU Media Access Unit for Token Ring networks

Mbps Millions of bits per second

MVS IBM's Multiple Virtual Storage operating system for mainframes

NFS Sun's Network File System

NetView IBM's network management system

OSF The Open Software Foundation

OSI OSI's Open Systems Interconnection protocol reference model

PSI Performance Systems International, Columbia's Internet service provider

PU An SNA Physical Unit

RJE Remote Job Entry protocol

RSCS IBM's Remote Spooling Communication Subsystem

Rlogin A UNIX-specific virtual terminal protocol

SAA IBM's Systems Application Architecture

SDLC IBM's Synchronous Data Link Control

SMTP The TCP/IP Simple Mail Transport Protocol

SNA IBM's Systems Network Architecture

SNMP The TCP/IP Simple Network Management Protocol

SQL Structure Query Language

Spectrum Cabletron's network management system

T1 A physical connection method, usually leased from the phone company, providing service at 1.544 million bits per second

TCP Transmission Control Protocol, the transport layer of TCP/IP

TCP/IP Transmission Control Protocol / Internet Protocol, the standard, open protocol used by the worldwide Internet

TELNET The TCP/IP virtual terminal protocol

TN3270 The TCP/IP virtual terminal protocol for 3270 emulation

Token Ring
A local area network technology in which stations communicate with each other at 4 or 10 Mbps over a shared cable in ring topology

TRN Abbreviation for Token Ring Network

TSO IBM's Time Sharing Option terminal monitor for MVS

UNIX A popular operating system developed at AT&T Bell Laboratories and noted for its portability

UTP Unshielded Twisted Pair wiring, such as our Rolm building wiring

UUCP UNIX-to-UNIX Copy Program

VM Virtual Memory, the name of the base operating system of our IBM mainframe

VT (1) The ISO Virtual Terminal protocol; (2) a DEC video terminal

WAN Wide Area Network

X The MIT X Window system

X.25 The CCITT network protocol

X.400 The CCITT electronic messaging protocol

X.500 The CCITT directory service

X/Open A consortium of corporations promoting open network protocols, etc.

Appendix IV. Attachments

1. 3270 Device Summary
2. CLIO Terminal Summary
3. Columbia Presbyterian Medical Center Network Summary
4. Current Morningside Academic Ethernet Backbone Map
5. Future Morningside Fiber Optic Network Backbone Map
6. Report of the Columbia University Network Architecture Task Force, September 1990

Table of Contents