Proceedings of the Seventeenth Internet Engineering Task Force

Pittsburgh Supercomputer Center
May 1-4, 1990

Corporation
for
National
Research
Initiatives

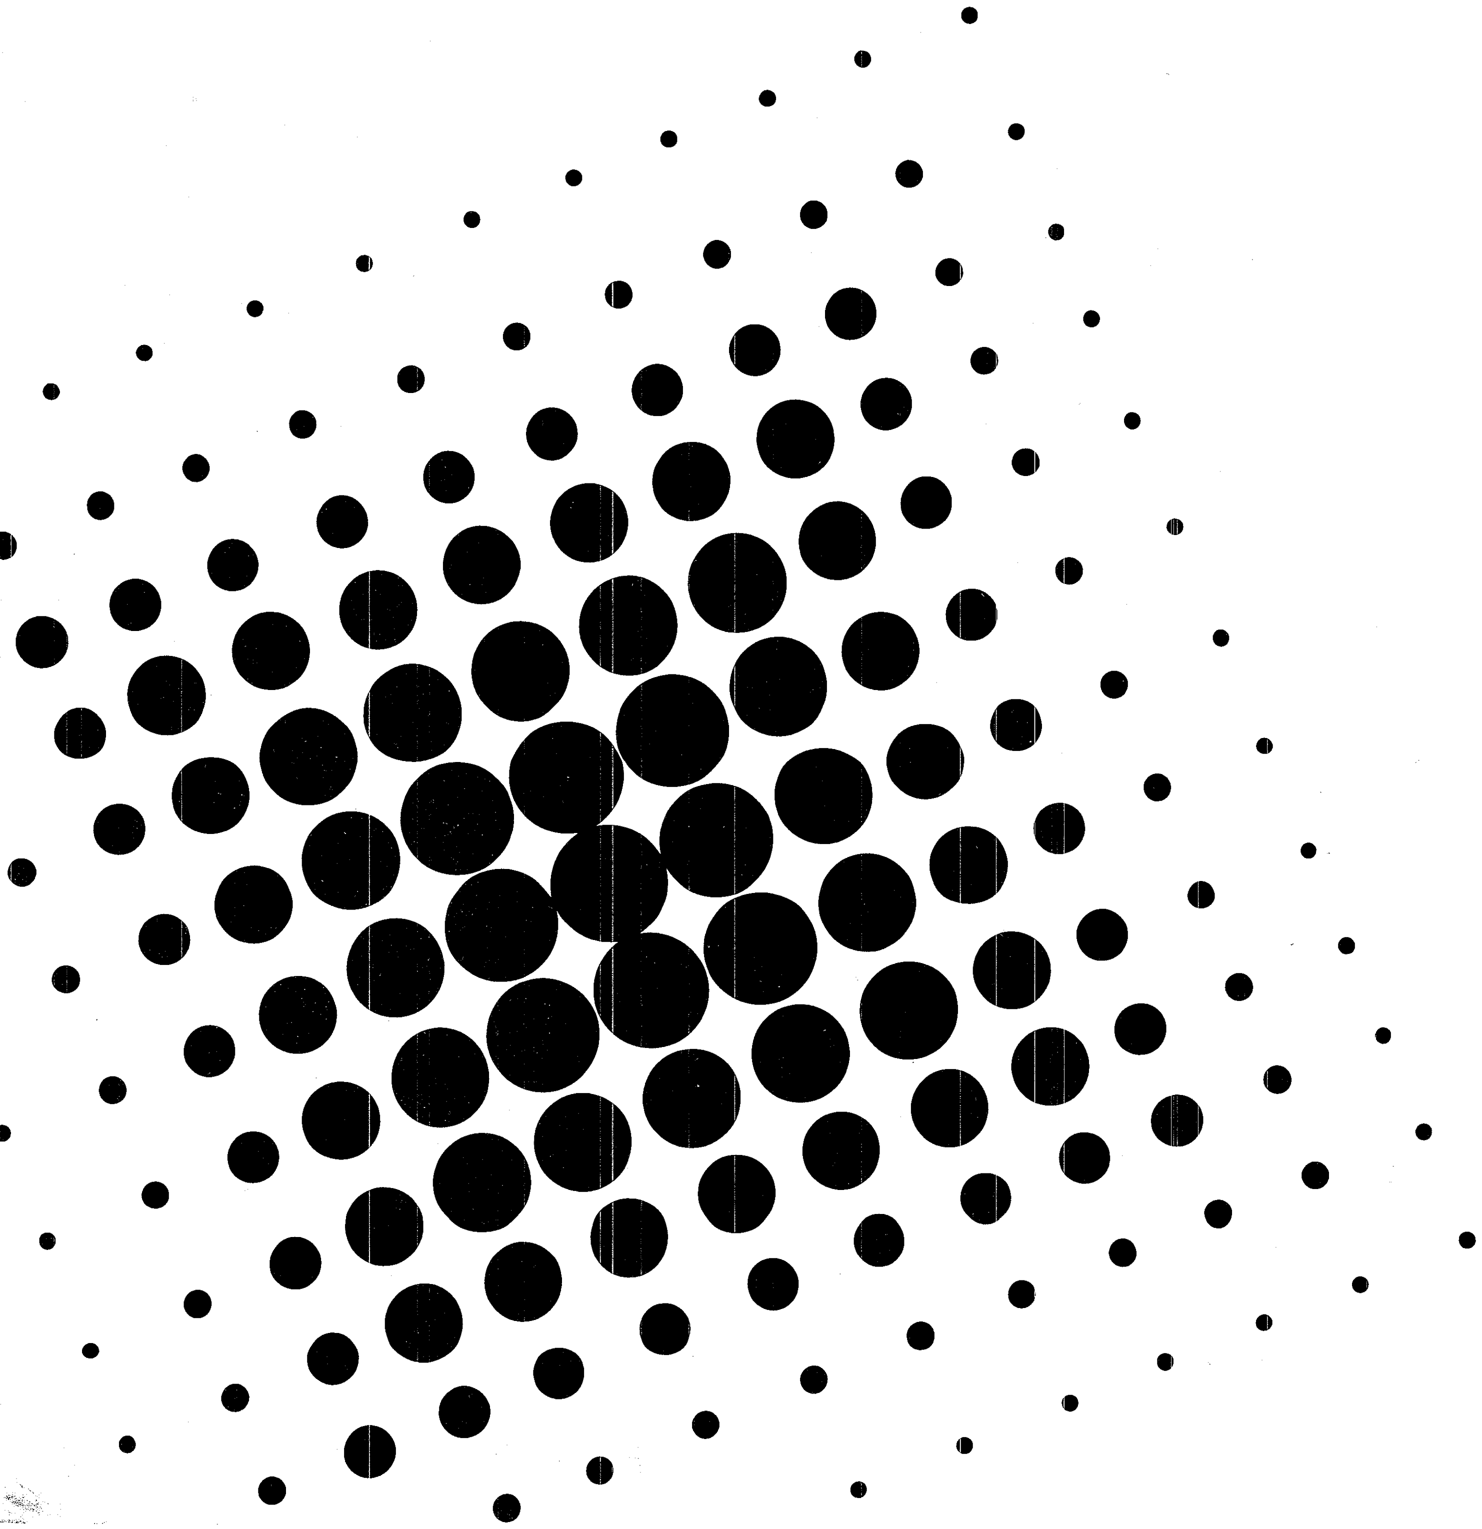

PROCEEDINGS OF THE SEVENTEENTH INTERNET ENGINEERING TASK FORCE

PITTSBURGH SUPERCOMPUTER CENTER
May 1-4, 1990

Compiled and Edited by
Phillip G. Gross
Gregory M. Vaudreuil

Corporation for National Research Initiatives
1895 Preston White Drive, Suite 100
Reston, Virginia 22091

This document and the IETF plenary it reports have been jointly sponsored by the National Science Foundation, the Defense Advanced Research Projects Administration, the Department of Energy, and the National Aeronautics and Space Administration through a grant from the the National Science Foundation (NCR-8820945).

ACKNOWLEDGEMENTS

The Seventeenth IETF plenary at the Pittsburgh Supercomputing Center held from May 1st - May 4th was a great success. On behalf of the Internet Engineering Task Force I would like to extend my thanks to Gene Hastings and the staff of the Pittsburgh Supercomputing Center for a very productive meeting. Thank you also to Richard Pethia of the CERT at the Software Engineering Institute for providing last minute meeting space as the attendee list grew. With over 240 attendees and over 45 working group sessions the meeting became quite a production.

Thank you to Bell of Pennsylvania, Prepnet, Carnegie Mellon University, and the Software Engineering Institute, and the many people who worked together to provided the excellent terminal facilities for the meeting attendees, especially the terminal cluster in the Holiday Inn hotel. With those efforts you have established a new standard for service to the community. What a Joy!

I would like to acknowledge the technical speakers, Allan Cargille, Claudia Topolcic, Philip Lehman, Ralph Droms, James Galvin, Yakov Rekhter, Tom Bajzek and Walt Burmeister, Tony Villasnor, and Phill Gross. These presentations were very informative. The network status presentations were given by Tony Hain, Milo Medin, Hans-Werner Braun, and Zbigniew Opalka.

And finally, a grateful acknowledgemnet and farewell to Monica Hart, who was invaluable in putting out these and past Proceedings. We found that it took us longer to put out the Proceedings without Monica's help. One positive result, however, is that this document now covers more than just the last IETF plenary. It is completely up-to-date through the middle of July.

Greg Vaudreuil / CNRI

Contents

Chairman's Message

Pittsburgh IETF

I would like to express my personal thanks to SEI, PSC, CMU, Prepnet, and Bell of Pennsylvania for hosting the May 1990 IETF meeting, and for setting a new standard in Internet connectivity. This is certainly a style to which we would like to become accustomed!

This meeting was attended by almost 250 persons. Thirty-three Working Groups (out of the total forty active IETF working groups) met in almost forty-five separate sessions. These numbers represented new high points for the IETF, and motivated us to take a closer look at IETF growth, activity and progress.

CNRI is developing a database facility for tracking all online IETF information. When completed, we will be able to provide the capability of querying for general information on IETF (including logistics information for upcoming meetings), and for specific information about Working Groups (including Working Group objectives, projected dates for accomplishing objectives, meeting minutes and Internet-Drafts).

We will use this locally to help us track IETF Working Group activity, but eventually we would also like to make this facility available as an anonymous TELNET service as a convenient way for interested parties to obtain information about IETF. The information below about IETF activities was derived using tools and data from the database.

Thanks go to Greg Vaudreuil (CNRI) for developing the database tools. It is our goal that most of the information now available in the quarterly IETF Proceedings or in the current online directories will eventually be available through the database tools and reports.

IETF Growth since January 1989

The following information about IETF attendance and growth is from the database.

Attendance for the last six IETF meetings:

12th IETF	Jan 1989	University of Texas	121
13th IETF	Apr 1989	Kennedy Space Center	112
14th IETF	Jul 1989	Stanford University	215
15th IETF	Oct 1989	University of Hawaii	138
16th IETF	Feb 1990	Florida State Univ.	191
17th IETF	May 1990	PSC/SEI/CMU	243

With the exception of the Stanford meeting (which may have been overly large because of proximity to computer industry), the above figures show a steady growth from around 100 to over 200 in the last year. Total attendance at these six meetings represents attendance by 500 different persons from 166 different organizations.

Repeat attendance by individal participants reveals a dedicated core of key IETF contributors. Twenty-three individuals have attended the past 6 meetings. Twenty-eight have attended 5 meetings, while another Forty-Six have attended at least 4. Nearly 100 folks have attended at least 2/3 of the recent meetings, an impressive statistic when attendance was only just over 100 for 3 of the those meetings.

When grouped by categories, we found that approximately 1/3 of the attendees were from vendors, 1/3 from government (DoD and civilian agencies), and more than 1/4 from universities and regional network operators.

IETF Activity and Progress since January 1989

Perhaps a more important measure of IETF activity is the number of active Working Groups and the number of RFCs produced over the same period. The following list shows the total number of Working Groups and the number which actually met at each meeting:

Date	Location	Total WGs	WG's Met
- Jan 1989	University of Texas	12	12
- Apr 1989	Kennedy Space Center	19	17
- Jul 1989	Stanford University	20	18
- Oct 1989	University of Hawaii	19	18
- Feb 1990	Florida State Univ.	38	32
- May 1990	PSC/SEI/CMU	40	33
- current (for UBC)		45 (approx)	

Notice that the number of Working Groups has shown a sharp increase since the creation of the IESG last fall. Following the first IESG meeting at the University of Hawaii, the number of Working Groups doubled.

During this general period, there were over 80 RFCs published relating to Internet technical activities. Of those RFCs, around 30 pertained to Internet standards. The IETF accounted for almost 30 percent of the total RFCs published and for 55 percent of all RFCs pertaining to standards. The IAB itself, together with the IRTF, accounted for almost another 30 percent, meaning that the IAB as an organization (i.e., including IETF and IRTF) accounted for almost 60 percent of all RFCs published in this period.

A version of this information will be presented and discussed at the UBC IETF meeting.

Wishing a Speedy Recovery to Gene Hastings

I was very saddened to hear that Gene Hastings, the host for our May meeting at PSC, suffered a very tragic event. Gene's house burned in early June. It was essentially a total loss. Gene, himself, suffered significant burns. Gene is the chair of the Network Joint Monitoring working group, and has been active in IETF since its earliest days.

During an FEPG meeting on June 12th, we called Gene in the hospital to wish him well. I was impressed by the strength of his spirit in the face of such tragedy and obvious pain. Gene said he had many contacts and offers of help from his friends and acquaintances from 15 years of working in the area of computer networking. That emphasized to me the very human side of computer networking in a poignant way.

Please join me in wishing Gene a speedy recovery. Although it is not likely that he will be at the August IETF meeting at UCB (and I will be very pleased if he proves me wrong!), he assured me that he would be active in the IETF as soon as possible. I know he was back on the network while still in the hospital via dial-in SLIP. Gene, we can't wait to have you back!

Final Agenda of the Sixteenth IETF

(May 1-4, 1990)

TUESDAY, May 1

9:00-12:00 am Morning Working Group Sessions

- Internet Security Policy (Richard Pethia/ Cert)
- Distributed File Systems (Peter Honeyman/ U of Michigan)
- Router Requirements (Philip Almquist/ Stanford, Jim Forster/cisco)
- User Connectivity (Kent England/ BBN)
- Open Routing (Martha Steenstrup/ BBN, Marianne Lepp/ BBN)
- Switched Megabit Data Service (George Clapp/ Ameritec and Mike Fidler/ Ohio State)
- Connection IP (Claudio Topolcic/ BBN)

1:00-4:00 pm Afternoon Working Group Sessions

- Network Printing (Leo McLaughlin/ Wollongong)
- User Documents (Tracy Laquey/ U-Texas, Karen Roubicek/ BBN)
- Router Requirements (Philip Almquist/ Stanford, Jim Forster/cisco)
- Switched Megabit Data Service (George Clapp/ Ameritec and Mike Fidler/ Ohio State)
- Open Routing (Martha Steenstrup/ BBN, Marianne Lepp/ BBN)
- SNMP Authentication (Jeff Schiller/ MIT)
- Connection IP (Claudio Topolcic/ BBN)
- OSI X.400 (Rob Hagens/ U-Wisc)

4:15-5:30 pm Technical Presentations

- "Bringing X.400 to the Internet" (Allan Cargille/ U-Wisc)
- "The ST 2 Protocol" (Claudio Topolcic/ BBN)
- "The National Andrew File System" (Philip Lehman/ Transarc)

WEDNESDAY, May 2

9:00-12:00 am Morning Working Group Sessions

- IP over Appletalk (John Veizades/ Apple)
- Site Security Policy Handbook (Paul Holbrook/ CERT) and Joyce Reynolds/ ISI)
- Multicast OSPF (Steve Deering/ Stanford)
- Topology Engineering (Scott Brim/ Cornell)
- Open Routing (Martha Steenstrup/ BBN, Marianne Lepp/ BBN)
- OSI Internet Management (Brian Handspicker/ DEC) and Lee LaBarre/ Mitre)
- Connection IP (Claudio Topolcic/ BBN)
- Alert Management (Louis Steinberg/ IBM)

1:00-4:00 pm Afternoon Working Group Sessions

- Transmission MIB/ FDDI MIB Joint Meeting (John Cook/Chipcom, Jeff Case/ U-Tenn)
- Management Services Interface Working Group (Oscar Newkerk/DEC)
- Network Joint Management (Gene Hastings/ PSC)
- Interconnectivity Working Group (Guy Almes/ Rice)
- Router Discovery and MTU Discovery (Jeff Mogul and Steve Deering)
- Switched Megabit Data Service (George Clapp/ Ameritec and Mike Fidler/ Ohio State)
- Open Routing (Martha Steenstrup/ BBN, Marianne Lepp/ BBN)
- Telnet (Dave Borman/ Cray)
- OSI NSAP Assignment (Richard Colella/ NIST)
- Transmission MIB (John Cook/ Chipcom)

4:15-5:30 pm Network Status Briefings

- "Energy Sciences Network Report" (Tony Hain)
- "Nasa Sciences Internet Report" (Milo Medin /NASA)
- "NSFnet Report" (Hans-Werner Braun/ MERIT)
- "Mailbridge Report" (Zbigniew Opalka/ BBN)

7:00- Evening Working Group Sessions

- Network Information Services Infrastructure (NISI) (Dana Sitzler/ MERIT)

THURSDAY, May 3

9:00-12:00 am	Morning Working Group Sessions

- Dynamic Host Configuration (Ralph Droms/ Bucknell)
- User Services Working Group (Joyce Reynolds/ ISI)
- Interconnectivity Working Group (Guy Almes/ Rice)
- Internet Accounting (Cindi Mills/ BBN)
- Point to Point Protocol Extentions (Steve Knowles/ FTP)
- Connection IP (Claudio Topolcic/ BBN)
- Decnet Phase IV MIB (Jon Saperia/ DEC)
- Domain Name System (Paul Mockapetris /ISI)
- OSPF Experience and Discussion B.O.F.(Rob Coltun/ UMD)

1:00-4:15 pm	Technical Presentations

- "The Knowbot Information Service" Ralph Droms / Bucknell
- "Privacy Enhanced Mail" James Galvin / TIS
- "The Border Gateway Protocol" Yakov Rekhter / IBM
- "Prepnet" Tom Bajzek / PREPNET and Walt Burmeister/ Bell of Pennsylvania
- "The FRICC/FNC, NREN, CCIRN" Tony Villasnor/ NASA
- "The FRICC/FNC Engineering Planning Group" Phill Gross/ NRI

4:30-7:00 pm	IETF Steering Group and Open Plenary Meeting

FRIDAY, May 4

9:00-11:30 am Working Group Area and Selected Working Group Presentations

- Applications Area (Russ Hobby/ UC Davis)
- Host and User Services Area (Phill Gross/ NRI and Joyce Reynolds/ ISI)
- Internet Services Area (Noel Chiappa/ Consultant-Proteon)
- Network Management Area (Dave Crocker/ DEC)
- Operations Area (Interim - Phill Gross/ NRI)
- OSI Interoperability Area (Ross Callon/ DEC and Rob Hagens/ U-Wisc)
- Routing Area (Bob Hinden/ BBN)
- Security Area (Steve Crocker/ TIS)

11:30-12:00 am Concluding Remarks (Phill Gross, NRI)

12:15 pm Adjourn

Chapter 1

IETF Overview

The Internet Engineering Task Force (IETF) has grown into a large open community of network designers, operators, vendors, and researchers concerned with evolution of the Internet protocol architecture and the smooth operation of the Internet. The IETF began in January 1986 as a forum for technical coordination by contractors working on the ARPANET, DDN, and the Internet core gateway system.

The IETF mission includes:

- Specifying the short and mid-term Internet protocols and architecture for the Internet,
- Making recommendations regarding Internet protocol standards for IAB approval,
- Identifying and proposing solutions to pressing operational and technical problems in the Internet,
- Facilitating technology transfer from the Internet Research Task Force, and
- Providing a forum for the exchange of information within the Internet community between vendors, users, researchers, agency contractors, and network managers.

Technical activity on any specific topic in the IETF is addressed within working groups. All working groups are organized roughly by function into eight technical areas. Each is led by an area director who has primary responsibility for that one area of IETF activity. These eight technical directors with the chair of the IETF compose the Internet Engineering Steering Group (IESG).

The current areas and directors, which compose the IESG, are:

IETF and IESG Chair:	Phill Gross/ NRI
1 Applications:	Russ Hobby/ UC-Davis
2 Host and User Services:	Craig Partridge/ BBN
3. Internet Services:	Noel Chiappa/ Consultant
4 Routing:	Robert Hinden/ BBN
5 Network Management:	Dave Crocker/ DEC
6. OSI Integration:	Rob Hagens/ U-Wisc and
7.	Ross Callon/ DEC
8. Operations:	Phill Gross/ NRI (interim)
9 Security:	Steve Crocker/ TIS
IESG Secretary:	Greg Vaudreuil/ NRI

The working groups conduct business during plenary meetings of the IETF, during meetings outside of the IETF, and via electronic mail on mailing lists established for each group. The IETF holds quarterly plenary sessions composed of working group sessions, technical presentations and network status briefings. The meeting are currently three and one half days long and includes an open IESG meeting.

Meeting reports, charters (which include the working group mailing lists), and general information on current IETF activities are available on-line for anonymous FTP from several Internet hosts including nnsc.nsf.net.

Information and logistics about upcoming meetings of the IETF are distributed on the IETF mailing list. To join the list or for general inquiries about the IETF, send a request to `ietf-request@isi.edu`.

1.1 Future IETF Meeting Sites

Summer 1990

University of British Columbia
Host: John Demco
July 31- August 3, 1990

Fall/Winter 1990

National Center for Atmospheric Research (NCAR)
The University of Colorado
Host: Don Morris and Carol Ward
December 4-7, 1990

Spring 1991

Washington University in St. Louis
Host: Guru Parulkar
March 11-14, 1991

1.2 On Line IETF Information

The Internet Engineering Task Force maintains up-to-date on-line information on all its activities. There is a directory containing Internet-Draft documents and a directory containing IETF working group information. All this information is available for public access at several locations. (See section 1.2.3)

The "IETF" directory contains a general description of the IETF, summaries of ongoing working group activities and provides information on past and upcoming meetings. The directory generally reflects information contained in the most recent IETF Proceedings and Working Group Reports.

The "Internet-Drafts" directory has been installed to make available, for review and comment, draft documents that will be submitted ultimately to the IAB and the RFC Editor to be considered for publishing as an RFC. Comments are welcome and should be addressed to the responsible person whose name and email addresses are listed on the first page of the respective draft.

1.2.1 The IETF Directory

Below is a list of the files available in the IETF directory and a short synopsis of what each file contains.

Files prefixed with a 0 contain information about upcoming meetings. Files prefixed with a 1 contain general information about the IETF, the working groups, and the internet-drafts.

FILE NAME	
0mtg-agenda	the current agenda for the upcoming quarterly IETF plenary, which contains what Working Groups will be meeting and at what times, and the technical presentations and network status reports to be given.
0mtg-logistics	the announcement for the upcoming quarterly IETF plenary, which contains specific information on the date/location of the meeting, hotel/airline arrangements, meeting site accommodations and travel directions.
0mtg-rsvp	a standardized RSVP form to be used to notify the support staff of your plans to attend the upcoming IETF meeting.
0mtg-schedule	current and future meeting dates and sites for IETF plenaries.
1id-abstracts	the internet drafts current on-line in the internet-drafts directory.
1id-guidelines	instructions for authors of internet drafts.
1ietf-overview	a short description of the IETF, the IESG and how to participate.
1wg-summary	a listing of all current Working Groups, the working group chairmen and their email addresses, working group mailing list addresses, and, where applicable, documentation produced. This file also contains the standard acronym for the working groups by which the IETF and Internet-Drafts directories are keyed.

Finally, Working Groups have individual files dedicated to their particular activities which contain their respective Charters and Meeting Reports. Each Working Group file is named in this fashion:

<standard wg abbreviation>-charter.txt

<standard wg abbreviation>-minutes-date.txt

The "dir" or "ls" command will permit you to review what Working Group files are available and the specific naming scheme to use for a successful anonymous ftp action.

1.2.2 The Internet-Drafts Directory

The Internet-Drafts directory contains the current working documents of the IETF. These documents are indexed in the file 1id-abstracts.txt in the Internet-Drafts directory.

The documents are named according to the following conventions. If the document was generated in an IETF working group, the filename is:

draft-ietf-<std wg abrev>-<docname>-<rev>.txt , or .ps

where <std wg abrev> is the working group acronym, <docname> is a very short name, and <rev> is the revision number.

If the document was submitted for comment by a non-ietf group or author, the filename is:

draft-<org>-<author>-<docname>-<rev>.txt, or .ps

where <org> is the organization sponsoring the work and <author> is the author's name.

For more information on writing and installing an Internet-Draft, see the file 1id-guidelines, "Guidelines to Authors of Internet-Drafts".

1.2.3 Directory Locations

The directories are maintained primarily at the NSFnet Service Center (NNSC). There are several "shadow" machines which contain the IETF and INTERNET-DRAFTS directories. These machines may be more convenient that nnsc.nsf.nsf.

To access these directories, use FTP. After establishing a connection, Login with username ANONYMOUS and password GUEST. When logged in, change to the directory of your choice with the following commands:

```
cd internet-drafts
cd ietf
```

Individual files can then be retrieved using the GET command:

```
get <remote filename>   <local filename>
e.g., get 00README      readme.my.copy
```

NSF Network Service Center Address: nnsc.nsf.net

The Defense Data Network NIC Address: nic.ddn.mil

Internet-drafts are also available by mail server from this machine. For more information mail a request:

To: service@nic.ddn.mil
Subject: Help

NIC staff are happy to assist users with any problems that they may encounter in the process of obtaining files by FTP or "SERVICE". For assistance, phone the NIC hotline at 1-800-235-3155 between 6 am and 5 pm Pacific time.

Pacific Rim Address: munnari.oz.au

The Internet-drafts on this machine are stored in Unix compressed form (.Z).

Europe Address: nic.nordu.net (192.36.148.17)

1.3 Guidelines to Authors of Internet Drafts

The Internet Drafts Directory is available to provide authors with the ability to distribute and solicit comments on documents they plan to submit as RFC's. Submissions to the Directory should be sent to **"internet-drafts@nri.reston.va.us"**. Unrevised documents placed in the Internet Drafts Directory have a maximum life of six months. After that time, they will either be submitted to the RFC editor or will be deleted. After a document becomes an RFC, it will be replaced in the Internet Drafts Directory with an announcement to that effect for an additional six months.

Internet Drafts (I-D's) are generally in the format of an RFC. This format is described in RFC 1111.

Following the practice of the RFCs, submissions are acceptable in postscript format, but we strongly encourage a submission of a matching ascii version (even if figures must be deleted) for readers without postscript printers and for online searches.

There are differences between the RFC and I-D format. The Internet Drafts are not RFC's and are not a numbered document series. The words "INTERNET-DRAFT" should appear in place of "RFC XXXX" in the upper left hand corner. The document should not refer to itself as a RFC or a Draft RFC.

The Internet Draft should not state nor imply that it is a proposed standard. To do so conflicts with the role of the IAB, the RFC editor and the IESG. The title of the document should not infer a status. Avoid the use of the terms Standard, Proposed, Draft, Experimental, Historical, Required, Recommended, Elective, or Restricted in the title of the draft. These are common words in the "Status of the Memo" section and may cause confusion if placed in the title.

The document should have an abstract section, containing a two-to-three paragraph description suitable for referencing, archiving, and announcing the document. The abstract should follow the Status of this Memo section. If the draft becomes an RFC, the Status of the Memo section will be filled in by the RFC editor with a status assigned by the IAB. As an Internet Draft, that section should contain a statement approximating one of the following statements:

1. This draft document will be submitted to the RFC editor as a standards document. Distribution of this memo is unlimited. Please send comments to ....
2. This draft document will be submitted to the RFC editor as an informational document. Distribution of this memo is unlimited. Please send comments to ....

If the draft is lengthy, please include on the second page a table of contents to make the document easier to reference.

1.4 IETF Working Group Summary (by Area)

Applications

Russ Hobby
`rdhobby@ucdavis.edu`

Domain Name System (dns)
Chairman: Philip Almquist `pvm@isi.edu`
WG mail: `namedroppers@nic.ddn.mil`
Status: continuing

Network Database (netdata)
Chairman: Clifford Lynch `lynch@postgres.berkeley.edu`
WG mail:
Status: new

Network FAX (netfax)
Chairman: Mark Needleman `mhn@stubbs.ucop.edu`
WG mail: `netfax@stubbs.ucop.edu`
Status: new

Network Printing Protocol (npp)
Chairman: Leo McLaughlin `ljm@twg.com`
WG mail: `print-wg@pluto.dss.com`
Status: continuing

TELNET (telnet)
Chairman: Dave Borman `dab@cray.com`
WG mail: `telnet-ietf@cray.com`
Status: continuing

Internet Draft: "Telnet Environment Option", 04/01/1990, Dave Borman
<draft-ietf-telnet-environment-00.txt>

Internet Draft: “Telnet Authentication Option”, 04/01/1990, Dave Borman <draft-ietf-telnet-authentication-00.txt>

Internet Draft: “Telnet Encryption Option”, 04/01/1990, Dave Borman <draft-ietf-telnet-encryption-00.txt>

Internet Draft: “Telnet Linemode Option”, 04/27/1990, Dave Borman <draft-ietf-telnet-linemodeoption-01.txt>

Internet Draft: “Telnet Data Compression Option”, 04/30/1990, Dave Borman <draft-ietf-telnet-compression-00.txt>

Host and User Services

Craig Partridge
craig@nnsc.nsf.net

Distributed File Systems (dfs)

Chairman: Peter Honeyman honey@citi.umich.edu
WG mail: dfs-wg@citi.umich.edu
Status: continuing

Dynamic Host Configuration (dhc)

Chairman: Ralph Droms droms@sol.bucknell.edu
WG mail: host-conf@sol.bucknell.edu
Status: continuing

Internet Draft: "Dynamic Configuration of Internet Hosts", 11/01/1989, Ralph Droms <draft-ietf-dhc-problem-stmt-00.txt and .ps>

Internet User Population (iup)

Chairman: Craig Partridge craig@nnsc.nsf.net
WG mail: ietf@venera.isi.edu
Status: continuing

Network Information Services Infrastructure (nisi)

Chairman: Dana Sitzler dds@merit.edu
WG mail: nisi@merit.edu
Status: new

Special Host Requirements (shr)

Chairman: Bob Stewart rlstewart@eng.xyplex.com
WG mail: ietf-hosts@nnsc.nsf.net
Status: new

User Connectivity (ucp)
Chairman: Dan Long `long@bbn.com`
WG mail: `ucp@nic.near.net`
Status: new

User Documents (userdoc)
Chairmen: Karen Roubicek `roubicek@nnsc.nsf.net`
Tracy LaQuey
WG mail: `user-doc@nnsc.nsf.net`
Status: continuing

Internet Draft: "Where to Start - A Bibliography of General Internetworking Information", 07/05/1990, K. Bowers, T. LaQuey,, J. Reynolds, K. Roubicek,, M. Stahl, A. Yuan <draft-ietf-userdoc-bibliography-00>

User Services (uswg)
Chairman: Joyce Reynolds `jkrey@venera.isi.edu`
WG mail: `us-wg@nnsc.nsf.net`
Status: continuing

Internet Services

Noel Chiappa
jnc@lcs.mit.edu

Connection IP (cip)
Chairman: Claudio Topolcic topolcic@bbn.com
WG mail: cip@bbn.com
Status: continuing

IP MTU Discovery (mtudisc)
Chairman: Jeff Mogul mogul@decwrl.dec.com
WG mail: mtudwg@decwrl.dec.com
Status: continuing

Internet Draft: "Path MTU Discovery", 07/05/1990, Jeff Mogul, S Deering <draft-ietf-mtudisc-pathmtu-01.txt>

IP over Appletalk (appleip)
Chairman: John Veizades veizades@apple.com
WG mail: apple-ip@apple.com
Status: new

IP over FDDI (fddi)
Chairman: Dave Katz dkatz@merit.edu
WG mail: FDDI@merit.edu
Status: continuing

Internet Draft: "A Proposed Standard for the Transmission of IP Datagrams over FDDI Networks", 05/05/1990, Dave Katz <draft-ietf-fddi-ipdatagrams-01.txt>

IP over Switched Megabit Data Service (smds)
Chairmen: George Clapp meritec!clapp@bellcore.bellcore.com
Mike Fidler ts0026@ohstvma.ircc.ohio-state.edu
WG mail: smds@nri.reston.va.us
Status: continuing

Internet Draft: "A Proposed Standard for the Transmission of IP Datagrams over SMDS", 07/18/1990, Joe Lawrence, Dave Piscitello <draft-ietf-smds-ipdatagrams-00.txt>

Point-to-Point Protocol Extentions (pppext)

Chairman: Stev Knowles `stev@ftp.com`
WG mail: `ietf-ppp@ucdavis.edu`
Status: continuing

Internet Draft: "The Point-to-Point Protocol (PPP) Initial Configuration Options", 04/11/1990, Drew Perkins <draft-ietf-ppp-options-03.txt>

Router Discovery (rdisc)

Chairman: Steve Deering `deering@pescadero.stanford.edu`
WG mail: `gw-discovery@gregorio.stanford.edu`
Status: continuing

Router Requirements (rreq)

Chairmen: Jim Forster `forster@cisco.com`
Philip Almquist `almquist@jessica.stanford.edu`
WG mail: `ietf-rreq@Jessica.Stanford.edu`
Status: continuing

Network Management

Dave Crocker
dcrocker@nsl.dec.com

Alert Management (alertman)
Chairman: Louis Steinberg louiss@ibm.com
WG mail: alert-man@merit.edu
Status: continuing

Internet Draft: "Managing Asynchronously Generated Alerts", 03/28/1990, Louis Steinberg <draft-ietf-alertman-asyncalertman-02.txt>

Bridge MIB (bridge)
Chairman: Fred Baker baker@vitalink.com
WG mail: bridge-mib@nsl.dec.com
Status: new

DECnet Phase IV MIB (decnetiv)
Chairman: Jon Saperia saperia%tcpjon@decwrl.dec.com
WG mail: phiv-mib@jove.pa.dec.com
Status: continuing

FDDI MIB (fddimib)
Chairman: Jeff Case case@utkux1.utk.edu
WG mail:
Status: new

Internet Accounting (acct)
Chairman: Cyndi Mills cmills@bbn.com
WG mail: accounting-wg@bbn.com
Status: new

LAN Manager (lanman)

Chairman: Jim Greuel jimg@cnd.hp.com
WG mail: lanmanwg@cnd.hp.com
Status: continuing

Internet Draft: "Management Information Base for LAN Manager Alerts", 06/30/1990, Jim Greuel, Amatzia BenArtzi <draft-ietf-lanman-alerts-00.txt>

Internet Draft: " Management Information Base for LAN Manager Management", 06/30/1990, Jim Greuel, Amatzia BenArzi <draft-ietf-lanman-mib-00.txt>

Management Services Interface (msi)

Chairmen: Oscar Newkerk newkerk@decwet.dec.com
Sudhanshu Verma verma@hpindbu.hp.com
WG mail: msiwg@decwrl.dec.com
Status: continuing

Internet Draft: "Management Services Interface", 07/13/1990, Oscar Newkerk <draft-ietf-msi-api-02.txt and .ps>

OSI Internet Management (oim)

Chairmen: Lee LaBarre cel@mbunix.mitre.org
Brian Handspicker bd@vines.dec.com
WG mail: oim@mbunix.mitre.org
Status: continuing

Internet Draft: "Tutorial on OSI Event Management, Alarm Reporting, and Log Control for TCP/IP Networks", 02/01/1990, Lee LaBarre <draft-ietf-oim-eventmanagement-00.txt and .ps>

Internet Draft: "OSI Internet Management: Management Information Base", 05/18/1990, Lee LaBarre <draft-ietf-oim-mib2-01.txt>

Internet Draft: "The Common Management Information Services and Protocols for the Internet (CMOT and CMIP)", 05/30/1990, U. Warrier, L. Besaw, B.D. Handspicker L. LaBarre <draft-ietf-oim-cmot-00.txt>

Remote LAN Monitoring (rlanmib)
Chairman: Mike Erlinger mike@mti.com
WG mail: rlanmib@decwrl.dec.com
Status: new

Simple Network Management Protocol (snmp)
Chairman: Marshall Rose mrose@psi.com
WG mail: snmp-wg@nisc.nyser.net
Status: continuing

Internet Draft: "Experimental Definitions of Managed Objects for the t1-carrier Interface Type", 04/23/1990, M. T. Rose, Fred Baker <draft-ietf-snmp-t1mib-00.txt>

Transmission Mib (transmib)
Chairman: John Cook cook@chipcom.com
WG mail: unknown
Status: continuing

OSI Integration

Ross Callon
callon@erlang.dec.com
Rob Hagens
hagens@cs.wisc.edu

Assignment of OSI NSAP Addresses (osinsap)

Chairman: Richard Colella   colella@osi3.ncsl.nist.gov
WG mail: ietf-osi-nsap@osi3.ncsl.nist.gov
Status: continuing

Internet Draft: "OSI NSAP Address Format For Use In The Internet", 07/10/1990, R Colella, R Callon <draft-ietf-osinsap-format-00>

OSI General (osigen)

Chairmen: Rob Hagens   hagens@cs.wisc.edu
Ross Callon   callon@erlang.dec.com
WG mail: ietf-osi@cs.wisc.edu
Status: continuing

OSI-X.400 (osix400)

Chairman: Rob Hagens   hagens@cs.wisc.edu
WG mail: ietf-osi@cs.wisc.edu
Status: continuing

Operations

Phill Gross (Interim)
pgross@nri.reston.va.us

Benchmarking Methodology (bmwg)
Chairman: Scott Bradner sob@harvard.harvard.edu
WG mail: bmwg@harvisr.harvard.edu
Status: continuing

Internet Draft: "Benchmarking Terminology", 07/13/1990, Scott Bradner <draft-ietf-bmwg-terms-00.txt>

Network Joint Management (njm)
Chairman: Gene Hastings hastings@psc.edu
WG mail: njm@merit.edu
Status: continuing

Topology Engineering (tewg)
Chairman: Scott Brim swb@devvax.tn.cornell.edu
WG mail: tewg@devvax.tn.cornell.edu
Status: continuing

Routing

Bob Hinden
hinden@bbn.com

ISIS for IP Internets (isis)
Chairman: Ross Callon callon@erlang.dec.com
WG mail: isis@merit.edu
Status: continuing

Internet Draft: "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", 01/01/1990, Ross Callon <draft-ietf-isis-spec-00.ps>

Interconnectivity (iwg)
Chairman: Guy Almes almes@rice.edu
WG mail: iwg@rice.edu
Status: continuing

Internet Draft: "Experimental Definitions of Managed Objects for the Border Gateway Protocol (Version 2)", 07/17/1990, Steven Willis, John Burruss <draft-ietf-iwg-bgp-mib-00.txt>

Multicast Extentions to OSPF (mospf)
Chairman: Steve Deering deering@pescadero.stanford.edu
WG mail: mospf@devvax.tn.cornell.edu
Status: continuing

Open Systems Routing (orwg)
Chairman: Martha Steenstrup msteenst@bbn.com
WG mail: open-rout-interest@bbn.com
Status: continuing

Internet Draft: "An Architecture for Inter-Domain Policy Routing", 02/20/1990, Marianne Lepp, Martha Steenstrup <draft-ietf-orwg-architecture-01.ps>

Private Data Network Routing (pdnrout)
Chairman: CH Rokitansky `roki@isi.edu`
WG mail: `pdn-wg@bbn.com`
Status: continuing

Internet Draft: "Assignment/Reservation of Internet Network Numbers for the PDN-Cluster", 06/01/1989, Carl-Herbert Rokitansky <draft-ietf-pdn-pdnclusternetassignm-00.txt>

Internet Draft: "Application of the Cluster Addressing Scheme to X.25 Public Data Networks", 08/01/1989, Carl-Herbert Rokitansky <draft-ietf-pdn-pdncluster-00.txt>

Internet Draft: "Internet Cluster Addressing Scheme", 11/01/1989, Carl-Herbert Rokitansky <draft-ietf-pdn-clusterscheme-00.txt>

Internet Draft: "X.121 Address Resolution for IP Datagram Transmission Over X.25 Networks", 04/23/1990, Carl-Herbert Rokitansky <draft-ietf-pdn-xarp-00.txt-00.txt>

Security

Steve Crocker
crocker@tis.com

IP Authentication (ipauth)
Chairman: Jeff Schiller jis@athena.mit.edu
WG mail: awg@bitsy.mit.edu
Status: continuing

Internet Security Policy (spwg)
Chairman: Richard Pethia rdp@sei.cmu.edu
WG mail: spwg@nri.reston.va.us
Status: continuing

SNMP Authentication (snmpauth)
Chairman: Jeff Schiller jis@athena.mit.edu
WG mail: awg@bitsy.mit.edu
Status: continuing

Internet Draft: "Administration of SNMP Communities", 07/05/1990, James Davin, James Galvin, Keith McCloghrie <draft-ietf-snmpauth-communities-01.txt>

Internet Draft: "Authentication and Privacy in the SNMP", 07/05/1990, James Galvin, Keith McCloghrie, James Davin <draft-ietf-snmpauth-authsnmp-02.txt>

Internet Draft: "Experimental Definitions of Managed Objects for Administration of SNMPCommunities", 07/05/1990, Keith McCloghrie, James Davin, James Galvin <draft-ietf-snmpauth-manageobject-02.txt>

Site Security Policy Handbook (ssphwg)
Chairmen: Paul Holbrook
Joyce Reynolds jkrey@venera.isi.edu
WG mail: ssphwg@cert.sei.cmu.edu
Status: new

1.5 Current Internet Drafts

This summary sheet provides a short synopsis of each Internet Draft available within the "Internet-Drafts" Directory at the NIC and NNSC.

"Privacy Enhancement for Internet Electronic Mail: Part IV – Certificate Requests and Related Forms", B. Kaliski, 04/01/1990
<draft-rsadsi-kaliski-privacymailpartiv-00.txt>

> This RFC documents the procedures for interacting with RSA Data Security Inc. (RSADSI) as a certifying authority for Internet privacy-enhanced mail. These procedures include registering organizations, registering organizational notaries, requesting signatures on certificates, and requesting signatures on certificate revocation lists (CRLs). The document also publishes the top-level distinguished name and public key in RSADSI's hierarchy.
>
> We intend this document, with the exception of Appendix A, to be a reference for implementors of ancillary privacy-enhanced mail software; it is not at the appropriate level for users of that software. However, the contracts and forms in Appendix A are intended for users.

"An Interim Approach to use of Network Addresses", S.E. Kille, 01/31/1990
<draft-ucl-kille-networkaddresses-00.ps>

> The OSI Directory specifies an encoding of Presentation Address, Which utilizes OSI Network Addresses as defined in the OSI Network Layer Standards. The OSI Directory, and any OSI application utilizing the OSI Directory must be able to deal with these Network Addresses. Currently, most environments cannot cope with them. It is not reasonable or desirable for groups wishing to investigate and use OSI Applications in conjunction with with the OSI Directory to have to wait for the lower layers to sort out. This note is a proposal for mechanisms to utilize Network Addresses.
>
> This document specifies an addressing convention to be used in conjunction with other protocols.

"A String Encoding of Presentation Address", S.E. Kille, 01/31/1990
<draft-ucl-kille-presentationaddress-00.ps>

> There are a number of Environments where a simple string encoding of Presentation address is desirable. This specification defines such a representation.

"X,500 and Domains", S.E. Kille, 01/31/1990
<draft-ucl-kille-x500domains-00.ps>

This document considers X.500 in relation to Internet/UK Domains. A basic model of X.500 providing a higher level and more descriptive naming structure is proposed, which gives a range of new management and user facilities over and above those currently available.

"Working Implementation Agreements On Network Management Functions, Services and Protocols", Robert Aronoff, 05/24/1990 <draft-nist-nmsig-implagreements-00.txt>

This is the Working Document of the Network Management Special Interest Group (NMSIG) of the OSI Implementors Workshop (OIW). The OSI Internet Management (OIM) Working Group agreements on CMIS/CMIP reference this document.

"Managing Asynchronously Generated Alerts", Louis Steinberg, 03/28/1990 <draft-ietf-alertman-asyncalertman-02.txt>

This draft defines mechanisms to prevent a remotely managed entity from burdening a manager or network with an unexpected amount of network management information, and to ensure delivery of "important" information. The focus is on controlling the flow of asynchronously generated information, and not how the information is generated. Mechanisms for generating and controlling the generation of asynchronous information may involve protocol specific issues.

There are two understood mechanisms for transferring network management information from a managed entity to a manager; request-response driven polling, and the unsolicited sending of "alerts". Alerts are defined as any management information delivered to a manager that is not the result of a specific query. Advantages and disadvantages exist within each method. This draft discusses these in detail.

"The Authentication of Internet Datagrams", Jeff Schiller, 08/01/1989 <draft-ietf-auth-ipauthoption-00.txt>

This draft RFC describes a protocol and IP option to allow two communicating Internet hosts to authenticate datagrams that travel from one to the other. This authentication is limited to source, destination IP address pair. It is up to host-based mechanisms to provide authentication between separate processes running on the same IP host. The protocol will provide for "authentication" of the datagram, not concealment from third party observers. By authentication, I mean that an IP host receiving a datagram claiming to be from some other IP host will be able (if both hosts are set up to authenticate datagrams between each other) to determine if in fact the datagram is from the host claimed, and that it has not been altered in transit.

"Benchmarking Terminology", Scott Bradner, 07/13/1990
<draft-ietf-bmwg-terms-00.txt>

> This memo discusses and defines a number of terms that are used in describing performance benchmarking tests and the results of such tests.
>
> The terms defined in this memo will be used in additional memos to define specific benchmarking tests and the suggested format to be used in reporting the results of each of the tests.

"Dynamic Configuration of Internet Hosts", Ralph Droms, 11/01/1989
<draft-ietf-dhc-problem-stmt-00.txt and .ps>

> This is a working document written by the Dynamic Host Configuration Working Group of the IETF. This document will be submitted as an RFC on February 12. Please respond with comments to the host-conf@rutgers.edu mailing list before that date or at the February, 1990 IETF meeting.

"A Proposed Standard for the Transmission of IP Datagrams over FDDI Networks", Dave Katz, 05/05/1990
<draft-ietf-fddi-ipdatagrams-01.txt>

> The goal of this specification is to allow compatible and interoperable implementations for transmitting IP datagrams and ARP requests and replies over FDDI networks.

"Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", Ross Callon, 01/01/1990
<draft-ietf-isis-spec-00.ps>

> This internet draft specifies an integrated routing protocol, based on the OSI Intra-Domain IS-IS Routing Protocol, which may be used as an interior gateway protocol (IGP) to support TCP/IP as well as OSI. This allows a single routing protocol to be used to support pure IP environments, pure OSI environments and dual environments. This specification was developed by the IS-IS working group of the Internet Engineering Task Force. Comments should be sent to "is-ismerit.edu".

"Experimental Definitions of Managed Objects for the Border Gateway Protocol (Version 2)", Steven Willis, John Burruss, 07/17/1990
<draft-ietf-iwg-bgp-mib-00.txt>

> This memo defines an experimental portion of the Management Information Base (MIB) for use with the Border Gateway Protocol [9,10] in TCP/IP-based internets.

"Management Services Interface", Oscar Newkerk, 07/13/1990
<draft-ietf-msi-api-02.txt and .ps>

The Management Services API defines Application Programming Interfaces which provide a set of services for the management of the objects in a heterogeneous, multivendor distributed computing environment.

The Management Services API is designed to allow for the development of portable management applications. The Management Services API insulate management application developers from the details of the management protocol and from the transport services used to route the management directives to the managed objects. It provides facilities to manage both local and remote objects in a seamless fashion.

"Path MTU Discovery", Jeff Mogul, S Deering, 07/05/1990
<draft-ietf-mtudisc-pathmtu-01.txt>

This memo describes a technique for dynamically discovering the maximum transmission unit (MTU) of an arbitrary internet path. It specifies a small change to the way routers generate one type of ICMP message. For a path that passes through a router that has not been so changed, this technique might not discover the correct Path MTU, but it will always choose a Path MTU as accurate as, and in many cases more accurate than, the Path MTU that would be chosen by current practice.

"The Knowbot Information Service", Ralph Droms, 12/01/1989
<draft-nri-droms-kis-00.txt and .ps>

Within the metanetwork of networks that exchange electronic mail, there are many directory services that provide partial coverage of network users; that is, directories with information about some subset of a particular network's user population. Searching the collection of available directories is time-consuming and requires knowledge of each directory's user interface. Although X.500 is currently under study as a basis for an Internet-wide directory service, it is unlikely that a universal user registry will be in place in the near future. The Knowbot Information Service provides a uniform interface to heterogeneous directory services that simplifies the task of locating users in the combined network.

"IP Routing Between U.S. Government Agency Backbones and Other Networks", Scott Brim, 01/01/1990
<draft-fricc-brim-BackboneRouting-01.txt>

This is an overview of how the agency backbones route IP (Internet Protocol) packets at this time, with any generalizations that can be made and statements of their differences. Also included are recommendations from the agency backbones about how other networks that connect to them can best set up their inter-administration routing.

"OSI Connectionless Transport Services on top of the UDP: Version 1", C. Shue, W. Haggerty, K. Dobbins, 11/01/1989
<draft-osf-shue-osiudp-00.txt>

This draft proposes a method for offering the OSI connectionless transport service (CLTS) in TCP/IP-based Internets by defining a mapping of the CLTS onto the User Datagram Protocol (UDP). If this draft becomes a standard, hosts on the Internet that choose to implement OSI connectionless transport services on top of the UDP would be expected to adopt and implement the methods specified in this draft. UDP port 102 is reserved for hosts which implement this draft. Distribution of this memo is unlimited.

"Implementation Agreements for Transport Service Bridges", M.T. Rose, 01/01/1990
<draft-ietf-rose-tsbridge-00.txt>

This draft reports implementation experience when building transport service bridges for OSI applications.

"Tutorial on OSI Event Management, Alarm Reporting, and Log Control for TCP/IP Networks", Lee LaBarre, 02/01/1990
<draft-ietf-oim-eventmanagement-00.txt and .ps>

This draft provides a tutorial on OSI mechanisms for event management, alarm reporting, and log control in TCP/IP networks. The mechanisms are based on ISO Draft Proposals and are expected to align with agreements developed by the National Institute of Standards and Technology (NIST) and the Network Management Forum (NMF). Also included is a mechanism for incorporating event flow control as defined in the Internet. It is proposed that systems implementing OSI management protocols for TCP/IP networks [1] should include the mechanisms described in this draft.

"OSI Internet Management: Management Information Base", Lee LaBarre, 05/18/1990
<draft-ietf-oim-mib2-01.txt>

This draft defines the Management Information Base (MIB) for use with the OSI network management protocol in TCP/IP based internets. It formats the Management Information Base (MIB-II) in OSI templates and adds variables necessary for use with the OSI management protocol.

"The Common Management Information Services and Protocols for the Internet (CMOT and CMIP)", U. Warrier, L. Besaw, B.D. Handspicker L. LaBarre, 05/30/1990
<draft-ietf-oim-cmot-00.txt>

> This memo is the output of the OSI Internet Management working group. As directed by the IAB in RFC 1052, it addresses the need for a long-term network management system based on ISO CMIS/CMIP. This memo contains a set of protocol agreements for implementing a network management system based on these ISO Management standards. Now that CMIS/CMIP has been voted an International Standard (IS), it has become a stable basis for product development. This profile specifies how to apply CMIP to management of both IP-based and OSI-based Internet networks. Network management using ISO CMIP to manage IP-based networks will be refered to as "CMIP Over TCP/IP" (CMOT). Network management using ISO CMIP to manage OSI-based networks will be refered to as "CMIP". This memo specifies the protocol agreements necessary to implement CMIP and accompanying ISO protocols over OSI, TCP and UDP transport protocols.

"An Architecture for Inter-Domain Policy Routing", Marianne Lepp, Martha Steenstrup, 02/20/1990
<draft-ietf-orwg-architecture-01.ps>

> We present an architecture for policy routing among administrative domains within the Internet. The objective of inter-domain policy routing is to synthesize and maintain routes between source and destination administrative domains, providing user traffic with the requested service within the constraints stipulated by the administrative domains transited. The architecture is designed to accommodate an Internet with tens of thousands of administrative domains.

"OSI NSAP Address Format For Use In The Internet", R Colella, R Callon, 07/10/1990
<draft-ietf-osinsap-format-00>

> This document provides alignment with U.S. GOSIP Version 2. GOSIP Version 2 has undergone the required public review and comment period prior to becoming a Federal Information Processing Standard (FIPS). It will be published as a FIPS by the end of calendar year 1990.

"Gateway Congestion Control Policies", A.J. Mankin, K.K. Ramakrishnan, 07/06/1990
<draft-ietf-pcc-gwcc-01.txt>

The growth of network intensive Internet applications has made gateway congestion control a high priority. The IETF Performance and

Congestion Control Working Group surveyed and reviewed gateway congestion control and avoidance approaches in a series of meetings during 1988 and 1989. The purpose of this paper is to present our review of the congestion control approaches, as a way of encouraging new discussion and experimentation. Included in the survey are Source Quench, Random Drop, Congestion Indication (DEC Bit), and Fair Queueing. The task remains for Internet implementors to determine and agree on the most effective mechanisms for controlling gateway congestion.

"Assignment/Reservation of Internet Network Numbers for the PDN-Cluster", Carl-Herbert Rokitansky, 06/01/1989
<draft-ietf-pdn-pdnclusternetassignm-00.txt>

This document contains a proposal for the reservation of Internet network numbers for the PDN-cluster and the assignment of these PDN-cluster networks to all national X.25 public data networks (DNICs), which are worldwide already in operation.

"Application of the Cluster Addressing Scheme to X.25 Public Data Networks", Carl-Herbert Rokitansky, 08/01/1989
<draft-ietf-pdn-pdncluster-00.txt>

In this document, the application of the Internet cluster addressing scheme to the international system of X.25 Public Data Networks is discussed and a new concept of hierarchical VAN-gateway algorithms for worldwide network reachability information exchange is proposed.

"Internet Cluster Addressing Scheme", Carl-Herbert Rokitansky, 11/01/1989
<draft-ietf-pdn-clusterscheme-00.txt>

In this document, the new concept of an addressing scheme, similar, but inverse to the subnetting scheme, is proposed, in which a set of Internet networks is associated to an Internet cluster. This "Cluster Addressing Scheme" is of interest especially for wide-area networks, whose structure should be visible to the outside world for (global) routing decisions. In addition, the use of an address-mask (called "Cluster-Mask") for routing decisions within the cluster is discussed.

"X.121 Address Resolution for IP Datagram Transmission Over X.25 Networks", Carl-Herbert Rokitansky, 04/23/1990
<draft-ietf-pdn-xarp-00.txt-00.txt>

"The Point-to-Point Protocol (PPP): A Proposed Standard for the Transmission of Multi-Protocol Datagrams Over Point-to-Point Links", Drew

Perkins, 03/01/1990
<draft-ietf-ppp-multidatagrams-02.txt>

> The Point-to-Point Protocol (PPP) provides a method for transmitting datagrams over serial point-to-point links. PPP is composed of three parts:
>
> 1. A method for encapsulating datagrams over serial links.
> 2. An extensible Link Control Protocol (LCP).
> 3. A family of Network Control Protocols (NCP) for establishing and configuring different network-layer protocols.
>
> This document defines the encapsulation scheme, the basic LCP, and an NCP for establishing and configuring the Internet Protocol (IP) (called the IP Control Protocol, IPCP).
>
> The options and facilities used by the LCP and the IPCP are defined in separate documents. Control protocols for configuring and utilizing other network-layer protocols besides IP (e.g., DECNET, OSI) are expected to be developed as needed.

"The Point-to-Point Protocol (PPP) Initial Configuration Options", Drew Perkins, 04/11/1990
<draft-ietf-ppp-options-03.txt>
"A Proposed Standard for the Transmission of IP Datagrams over SMDS", Joe Lawrence, Dave Piscitello, 07/18/1990
<draft-ietf-smds-ipdatagrams-00.txt>

> This memo describes an initial use of IP and ARP in an SMDS environment configured as a logical IP subnet, LIS (described below). The encapsulation method used is described, as well as various service-specific issues. This memo does not preclude subsequent treatment of SMDS in configurations other than LIS; specifically, public or inter-company, inter-enterprise configurations may be treated differently and will be described in future documents.

"Experimental Definitions of Managed Objects for the t1-carrier Interface Type", M. T. Rose, Fred Baker, 04/23/1990
<draft-ietf-snmp-t1mib-00.txt>

"Administration of SNMP Communities", James Davin, James Galvin, Keith McCloghrie, 07/05/1990
<draft-ietf-snmpauth-communities-01.txt>

> Simple Network Management Protocol (SNMP) specification allows for the authentication of management operations by a variety of authentication algorithms. This memo defines two strategies for administering

SNMP communities based upon either the SNMP authentication algorithm or the SNMP authentication and privacy algorithm. Insofar as the administration of SNMP communities based upon the trivial authentication algorithm may be realized by straightforward application of familiar network management techniques, administration of such communities is not directly addressed in this memo.

"Authentication and Privacy in the SNMP", James Galvin, Keith McCloghrie, James Davin, 07/05/1990
<draft-ietf-snmpauth-authsnmp-02.txt>

The Simple Network Management Protocol (SNMP) specifica- tion allows for the authentication of network management operations by a variety of authentication algorithms. This memo specifies alternatives to the trivial authentication algo- rithm. It also describes an abstract Authentication Service Interface (ASI) by which SNMP-based management applications or agents may–in a convenient and uniform way–benefit from the algorithms described here and a wide range of others. The terms of the ASI are used to describe three distinct algorithms, including one with support for privacy.

"Experimental Definitions of Managed Objects for Administration of SNMPCommunities", Keith McCloghrie, James Davin, James Galvin, 07/05/1990
<draft-ietf-snmpauth-manageobject-02.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets. In particular, it describes a representation of the authentication communities defined in the companion memo: Authentication and Privacy in the SNMP as objects in the Internet Standard MIB. These definitions are consistent with the administrative strategies set forth in the companion memo: Administration of SNMP Communities.

"Telnet Linemode Option", Dave Borman, 04/27/1990
<draft-ietf-telnet-linemodeoption-01.txt>

Linemode Telnet is a way of doing terminal character processing on the client side of a Telnet connection. While in Linemode with editing enabled for the local side, network traffic is reduced to a couple of packets per command line, rather than a couple of packets per character typed. This is very useful for long delay networks, because the user has local response time while typing the command line, and only incurs the network delays after the command is typed. It is also useful to reduce costs on networks that charge on a per packet basis.

"Telnet Environment Option", Dave Borman, 04/01/1990 <draft-ietf-telnet-environment-00.txt>

"Telnet Authentication Option", Dave Borman, 04/01/1990 <draft-ietf-telnet-authentication-00.txt>

"Telnet Encryption Option", Dave Borman, 04/01/1990 <draft-ietf-telnet-encryption-00.txt>

"Telnet Data Compression Option", Dave Borman, 04/30/1990 <draft-ietf-telnet-compression-00.txt>

"Where to Start - A Bibliography of General Internetworking Information", K. Bowers, T. LaQuey,, J. Reynolds, K. Roubicek,, M. Stahl, A. Yuan, 07/05/1990 <draft-ietf-userdoc-bibliography-00>

The intent of this bibliography is to offer a representative collection of resources of information that will help the reader become familiar with the concepts of internetworking. It is meant to be a starting place for further research. There are references to other sources of information for those users wishing to pursue, in greater depth, the issues and complexities of the current networking environment.

Chapter 2

Steering Group Report

2.1 Minutes of the May 4th Meeting

The Internet Engineering Steering Group met during the open plenary on Thursday May 4th. The topics discussed were a status report on the evolving IAB standards process, and a report on IAB standards actions, and the ANSI initiative to standardize TCP/IP. Steve Crocker led a discussion on the important topic of network security.

2.1.1 The IAB Standards Process

The IAB is soliciting suggestions for replacing the practice of simply labeling standards RFC's with a one word "Requiredness Level". This system does not give enough information to be useful, and often leads to confusion.

The plenary had several views of using a requirement level. It was accepted that the "Requirement Level" does not have fine enough granularity. This view was well illustrated. What is necessary for a router is not necessary for a terminal server, and what is reasonable for a large multitasking system may not be for a PC.

The authority of the IAB is limited. Ultimately the decisions on which protocols are implemented and deployed in the internet is a result of user needs and the market. While the IAB may recommend a protocol, it is the users who must demand products, and vendors who must supply products. The recommendations may be useful to users in specifying products.

Many felt that the concept of a recommendation for use was very helpful. There are many persons and corporations who buy internet products who may not have the background or experience to specify specific protocols, but who appreciate general guidelines. This need is a motivating factor behind the Host Requirements document and the current Router Requirements effort.

While some argued that individuals who do not understand the protocols should not be writing requests for proposals, other argued that even for those who understand the technology, there is great utility in being able to justify purchasing requirement decisions by relying on the experience and authority of the IAB.

The accuracy of the requirements documents was an issue. The Host Requirements RFC's were an attempt at creating a complete guide to host usage of tcp-ip. It was a profile and an implementation guide in one document. By linking the two concepts into one document, the changing recommendations in the implementation guide are causing the document to be out of date, even as the protocol profiles remain relevant. It is now turning up to have shortfalls in specifying requirement for specialized hosts.

Vint Cerf, chair of the IAB, participated in these discussions. He, and the other IAB members present will relay the IETF discussion to the IAB.

2.1.2 IAB Standards Actions

Grandfathered Protocols

Document IESG recommendations	Protocol Name IAB Action
RFC 407 Historical	RJE - Remote Job Entry HISTORICAL
RFC 569 Historical	NETED - Network Standard Text Editor HISTORICAL
RFC 734 Draft Standard	SUPDUP Pending: Note (1)
RFC 742 Draft Standard	Finger Pending: Note (1)
RFC 818 Historical	RTELNET - Remote Telnet HISTORICAL
RFC 887 Historical	RLP - Resource Location Protocol Pending: Note (2)
RFC 913 Historical	SFTP - Simple File Transfer Protocol Pending: Note (2)
RFC 937 Historical	POP2 - Post Office Protocol, V. 2 HISTORICAL
RFC 953 Historical	Hostname Pending: Note (1)
RFC 954 Draft Standard	NICNAME - WhoIs DRAFT STANDARD
RFC-977 Proposed Standard	NNTP - Network News Transfer Protocol PROPOSED STANDARD
RFC 996 Historical	STATSRV - Statistics Server HISTORICAL

RFC 987, 1026 Experimental	Mapping Between X.400 and RFC 822 mail Pending: Note (3)
RFC 1037 (No IESG recommendation)	NFILE (Awaiting recommendation)
RFC 1045 Experimental	VMTP EXPERIMENTAL
RFC 1056 (No IESG recommendation)	PCMAIL (Awaiting recommendation)
RFC 1057 Proposed Standard	Sun Remote Procedure Call Information Only - Cerf will discuss with Sun
RFC 1058 Proposed Standard	Routing Information Protocol DRAFT STANDARD
RFC1081, 1082 (No IESG recommendation)	POP3 (No IAB action required)
RFC 1090 Experimental	SMTP over X.25 (No IAB action required)
RFC 1094 Proposed Standard	Sun Network File System Information Only - Cerf will discuss with Sun

Other Protocols:

RFC 1006 Draft Standard	ISO Transport on TCP DRAFT STANDARD
RFC 1098 Standard	SNMP - Simple Network Management Protocol STANDARD, Note (4)
RFC 1065 Standard	SMI - Structure for Managed Information STANDARD, Note (4)
RFC 1066 Standard	MIB 1 - Management Information Base STANDARD, Note (4)

RFC 1131 Proposed Standard	OSPF PROPOSED STANDARD
RFC-1134 Draft Standard	PPP - Point-to-Point Protocol Pending: Note (5)
Internet Draft Proposed Standard	PPP Initial Configuration Options Pending
Internet Draft Proposed Standard	MIB 2 PROPOSED STANDARD
Internet Draft Experimental	SNMP OSI MIB (No IAB action required)
Internet Draft Experimental	SNMP over OSI (No IAB action required)
Internet Draft Proposed Standard	An Interim Approach to Network Addresses [Kille] Pending: Note (6)
Internet Draft Proposed Standard	A String Encoding of Presentation Address [Kille] Pending: Note (6)
Internet Draft (No IESG recommendation)	BGP – Border Gateway Protocol (Awaiting recommendation)

Note 1: It has been suggested this remain a Proposed Standard, pending further review by IESG. This has been given low priority for IAB discussion, so no further action has taken place yet.

Note 2: It has been has suggested this become Experimental rather than Historical. This has been given low priority for IAB discussion, so no further action has taken place yet.

Note 3: This was only recently taken up by the IAB. One person has questioned whether moving it back to Experimental is desirable or appropriate. [The IESG recommendations later proved a miscommunication, and the status of the RFC was left unchanged]

Note 4: The IAB has agreed on the following Requirement Level and applicability statement for these protocols:

The SMI, MIB I, and SNMP are to be shown as Recommended in the IAB Official Protocol Standards document, with the additional note:

> The Internet Activities Board recommends that all IP and TCP implementations be network manageable. This implies implementation of the Internet MIB (RFC-1066) and at least one of the two recommended management protocols SNMP (RFC-1098) or CMOT (RFC-1095). It should be noted that SNMP is a full Internet standard while CMOT is a draft standard at this time.
>
> See also the Host and Router Requirements RFCs for more specific information on the applicability of this standard.

Note 5: The relationship to OSI for multiprotocol routers has been questioned.

Note 6: These are still under review by the IAB.

There was significant discussion on the issue of standardizing vendor proprietary protocols. There is precedent for the standardization of commercial protocols. Control of the Ethernet specification was given to IEEE by Xerox. In most cases, the standards organization has required change control over the protocol. Without change control, the standardization process can be manipulated by marketing concerns. Acceptance of protocols from other public, open organizations is less of a problem for this reason, because they are less likely to be manipulated by a single vendor.

For the Internet community, the standards issue becomes a bit complicated. There are internet protocols based on proprietary protocols. For example, the Paladium printing protocol is based on Sun Microsystems RPC. Many efforts in the IETF are closely coupled with vendor protocols, especially in the area of distributed file systems.

2.1.3 ANSI Initiative on Standardizing Internet Protocols

Vint Cerf presented an overview of the effort underway in ANSI to standardize the core TCP/IP protocols. (See the following slides) This process would enter the TCP-IP protocols into a suite of American National Standards. Further action on this matter now rests with ANSI X3S3.3. The IAB is interested on only in forwarding stable protocols for ANSI consideration. There are some protocols, such as the dual IS-IS where joint development is appropriate.

ANSI X3S3.3 (Ch. L. Chapin)
New Work Proposals

April 18, 1990 - TUCSON, AZ

1) IP, ICMP, TCP, UDP ⇒ ANS.

- Consolidate existing RFCs for each protocol and provide to X3S3.3 for balloting
- Any issues arising or new technical work would be remanded to IAB + IETF for resolution.

Tabled Apr 20. To be revisited ~June 29, 1990 X3S33 mtg.

L. Chapin has action on preparing work proposals in concert w/IAB.

IAB has tentatively approved.

2) Protocols outside of transport and network layers are outside X3S3.3 parview.

issue: should other elements of TCP/IP Suite be brought into ANS orbit?

3) IS-IS dual routing
+ OSPF also proposed for work by ANSI X3S3.3.
<u>Not resolved</u>. Revisit June 26.

4) BRP variant of BGP is an on-going item of work.

Security Area

Working Groups

- Authentication
  - SNMP
  - IP
- Security Policy
- new! Site Security Policy Handbook

Collateral Activities

- PSRG -> Privacy Enhanced Mail
- Telnet -> Auth, Integrity, Encryption
- Remote Printers -> Auth
- BGP -> Auth

Lost?

- IPSO

Emerging Topics

- General Purpose Authentication Protocol? Protocol Component?
- Security Analysis of Routing Protocols
- Operational Effectiveness
  - Security Testing Tools
  - Organizational Issues
- Legal & Liability Issues

Join Up!

GP Authentication

- Position papers on model
- Jim Galvin & Jeff Schiller
- Mail to crocker@tis.com
- Formal announcement RSN

Routing Protocol Analysis

- Security (Authentication, ...)
- Robustness
- Volunteers: crocker@tis.com

PSRG Meeting Vancouver

- PEM
  - Initial Distribution
  - Discussion
- Security Architecture Framework
- Joint meetings with IETF WGs

2.2 The IESG Standards Process

The following is a proposed standards process to be implemented by the IESG.

2.2.1 Internet-Drafts

Internet Drafts are posted at the request of the author with the following restrictions.

1. The Internet-Draft must conform to the Guidelines for Authors of Internet-Drafts". This includes strict enforcement of the Status of the Memo Section, and a one paragraph abstract.
2. The Internet-Draft will be announced to the IETF mailing list with an announcement derived from the abstract. In effect, the abstract is an announcement. The announcement should be one to two paragraphs, of preferably no more than 15 lines of text.
3. If the draft is a protocol specification, the "Status of the Memo" section should have the following words:

   > This document is not an internet standard. This draft document will be submitted to the RFC editor as a protocol specification. Distribution of this memo is unlimited. Please send comments to ....

   or if the draft is an informational document:

   > This draft document will be submitted to the RFC editor as an informational Document. Distribution of this memo is unlimited. Please send comments to ....

2.2.2 Internet-Drafts to RFC's

Standards Documents

1. The document author submits the Internet-Draft to the IESG by mailing a request to the IESG-Secretary@nri.reston.va.us.
2. The IESG Secretary announces the pending consideration of the Internet-Draft as a standard to the IAB, IESG, and IETF. This notice will include:
    - Timeframe for consideration, including the date of the plenary session it will be considered at.
    - Originating Working Group,
    - A brief abstract extracted from the Internet-Draft,
3. The authors of the document may be invited to give a technical presentation to the IETF plenary to describe the protocol and answer any questions that may arise.
4. The IESG may review the draft in open session at the next IETF plenary session. Items to be considered are:
    - Does this document meet the standards for a well defined specification?
    - Is this document considered implementable? A Proposed Internet Standard is preferred to have at least one implementation.
5. If there are significant concerns expressed, either technical or political, the IESG may at it's discretion:
    - Accept the draft,
    - Remand the draft back to the Working Group for further work,
    - Or submit the document for an independent technical review.
6. After all questions are resolved, the IESG formulates a recommendation to the IAB.
    - The document will be submitted by the IESG-secretary to the IAB via the RFC Editor CC'ed to the IETF list.
    - The submission will include a recommended "Status of the Memo" Section.

Informational Documents

- The document author submits the internet-draft to the IESG by mailing a request to the IESG-Secretary@nri.reston.va.us.
- The IESG-Secretary will consult with the IESG Area Director, CC'ing the IESG Mailing List, and if there are no objections, send the document to the RFC Editor.

2.2.3 Proposed Standard to Draft Standard

1. The document author submits the edited and updated Proposed Standard to the IESG by mailing a request to the IESG-Secretary@nri.reston.va.us. It would be useful for this document to include a section detailing changes from the Proposed Standard document, and any significant information useful to implementors.
2. The IESG Secretary will announce the consideration of the RFC for elevation to Draft Standard to the IAB, IESG, and IETF. This notice will include:
    - Timeframe for consideration, including the date of the plenary session it will be considered at.
    - A generic "Status of the Memo Section"
    - A brief abstract extracted from the RFC
    - A Pointer to the revised document in the Internet-Drafts directory if the document has been significantly revised.
    - And, send the revised document to the IAB.
3. The IESG reviews the RFC in open session at the next IETF plenary session. Items to discuss are:
    - Has the protocol been a Proposed Standard for at least 6 months?
    - Does this protocol meet the requirements as an independently implementable specification? This is evidenced by multiple independent interoperable implementations of the Proposed Standard as refined in the submitted Draft Standard RFC.
    - Is this protocol considered operationally stable? A Draft Standard is preferred to have significant operational experience.
    - If there are significant concerns expressed, either technical or political, the IESG may at it's discretion submit the document for an independent technical review, or remand the document to the Working Group for further work.
4. After all questions are resolved, the IESG formulates a recommendation.
    - The document will be submitted by the IESG-secretary to the IAB Via the RFC Editor CC'ed to the IETF list.
    - The submission will include a recommended "Status of the Memo" Section.

2.2.4 Draft Standard to Full Standard

1. The document author submits the edited Draft Standard to the IESG by mailing a request to the IESG-Secretary@nri.reston.va.us. It is recommended that this document include a section detailing any changes from the Draft Internet Standard document. This should include information necessary for operationally usage in the Internet. Implementation and operational experience may conveyed in a companion document.
2. The IESG Secretary will announce the consideration of the RFC for elevation to Full Standard to the IAB, IESG, and IETF. This notice will include:
   - Timeframe for consideration, including the date of the plenary session it will be considered at.
   - A brief abstract extracted from the RFC
3. The IESG will review the RFC in open session at the next IETF plenary session. Items to discuss are:
   - Has the protocol been a Draft Internet Standard for at least 6 months?
   - Does this protocol meet the requirements as an completely defined specification with multiple independent interoperable implementations of the Draft Standard RFC?
   - Does this protocol meet the requirement as an operationally stable Protocol as evidenced by widespread deployment and operational experience.
   - If there are significant concerns expressed, either technical or political, the IESG may at it's discretion submit the document for an independent technical review, or remand the document to the Working Group for further work.
4. After all questions are resolved, the IESG formulates a recommendation.
   - The document will be submitted by the IESG-secretary to the IAB Via the RFC Editor CC'ed to the IETF list.
   - The submission will include a recommended "Status of the Memo" Section.
   - Included in the recommendation to the IAB should be a short statement on status and consequences, an "Environmental Impact Rreport", on the cost and benefit of deploying the protocol.

Chapter 3

Area and Working Group Reports

3.1 Applications Area

Director: Russ Hobby/UC Davis

WORKING GROUPS ACTIVE AT PITTSBURGH

Domain Name System - This Working Group has a new chair, Phillip Almquist, who will determine if there are any areas on which the group needs to work. Anyone who thinks there are problem areas in the Domain Name System should contact Phillip.

Network FAX - There was a short meeting to define the direction of this Working Group. Mark Needleman will write a requirements document to reflect this discussion.

Network Printing Protocol - This group has produced a document defining LPR. Their main work was looking at Palladium, the printing protocol used in Project Athena at MIT. This protocol may meet the needs for an Internet printing protocol.

TELNET - the Working Group has produced two documents. The Linemode document is ready to be submitted to be a proposed standard as it has had implementations and the RFC has had a few minor changes. The Working Group has also produced a document for the Environment Option and is ready to summit it to be a proposed standard. Progress was made on new options defining authentication, encryption, and Tn3270.

WORKING GROUPS NOT MEETING AT PITTSBURGH

Network SQL - This is a new Working Group and will define a standard for the use of SQL databases over TCP/IP networks. It is viewed that the work that has been done to define SQL over OSI can be mapped into TCP/IP. The chair of this Working Group is Clifford Lynch (lynch@postgres.berkeley.edu) and those interested should contact him for information and to be added to the mail list.

APPLICATIONS ON THE INTERNET

We, as the engineers of the Internet, have had a tendency to look at the network from the bottom of the protocol stack looking up. We have mainly focused on how we get the bits across the network and not so much how they are used. We have now created a large network with many users. It is time for some of us to look at it from their point of view.

To begin, we need to answer a few questions. What do the users want to do with the network? What resources are available on the network today? Do they meet the needs and expectations of the users? If not, What do we do next?

Let's look at the network from the users view point. There are at least three, and probably more, types of applications. First there are applications for the searching, retrieval, and distribution of information. Next there are applications for personal communications and finally there are operational applications for use in the general computing environment.

INFORMATION APPLICATIONS

There is already a vast amount of information on the Internet, but it is not easy to find or access. We need applications to help us search for and retrieve information, plus standard formats for that information so that we can do something with it after we get it. There are many types of information that can be accessed by computer, but let's look at some that we have today but need to provide better access.

One information service is for information on people. There are whois servers on the Internet and new projects using X.500. Through X.500 and the White Pages Project the Internet has a good start of providing information of the Internet population and beyond. We just need to help the implementation of X.500 on the Internet.

Another use is library type functions. Currently many facilities have there card catalogs and other bibliographic information on line, but we are seeing more of the actual information itself on line. The biggest problem is how to find it on the network and once you do how do you get and use it.

Software sharing and distribution is a popular information sharing function. Many of use have seen the advantage of "anonymous FTP", but again, finding the software in the first place is not easy. For commercial producers of software, licensing of network distributed software (and other information as well) needs to be considered.

PERSONAL COMMUNICATIONS APPLICATIONS

Three types of personal communications that people seem to want are person-to-person, person-to-group, and calendar/scheduling. Person-to-person is communications to one person or a small know group of people and include functions such as electronic mail, talk/chat, video conference and, of course, the telephone.

Person-to-group are broadcast type of communications where you are communicating with a large unknown group. This includes services such as forums and bulletin boards. USENET is probably the most popular type of this service currently available.

For the third type, imagine that you could maintain your personal calendar on your own computer. Now imagine that your calendar can talk to all other calendars on computers all over the Internet and schedule people, rooms, and other resources. This is the type of functionality that people want of network calendar/scheduling.

OPERATIONAL APPLICATIONS

There are network functions that are associated more with distributed computing rather than communications. These generally have to do with resource sharing of devices such as printers, disks, backup storage and to some extent, CPUs. Back in the "old days" of computers and mainframes, users had to know where each peripheral was and how to access it. Operating systems have now made that invisible to the user on individual computers.

When it comes to network resources, however, we have jumped back twenty years. You still need to know where each resource is on the network and how to access it. We need to use what we have learned in operating systems and apply it to networks. Think of the network as a computer buss with lots of CPUs and peripherals hung on it. THE NETWORK IS THE COMPUTER. Now we just need to write a user friendly operating system for this computer.

The real problem, of course, is that with single computers, it has been just one vendor that has had to coordinate within itself. With networks, we are operating in a multi-vendor environment and coordination means that we need standards.

TOOLS TO BUILD APPLICATIONS

So now we have all these nice applications that we want to write, or make the old ones better. Many of them have lower level functions in common, such as authentication, remote procedure calls, remote file operations, and remote data bases. They also need to agree on formats for information, such as character sets, graphics format, file structures and command syntax. Most of these tools and formats do not have a standard definition for the Internet.

What do we do now? The current Working Groups are basing their work on the assumption that these tools will be there. The primary tools that seem to be needed now for these Working Groups are authentication and remote procedure calls, but the others will soon be needed too. One factor that adds to the confusion is the fact that other bodies are also trying to decide on standard tools and formats and failing to come to agreement. For the Internet, where interoperability has been the key to success we need to agree on a common direction. So, yes, what do we do now?

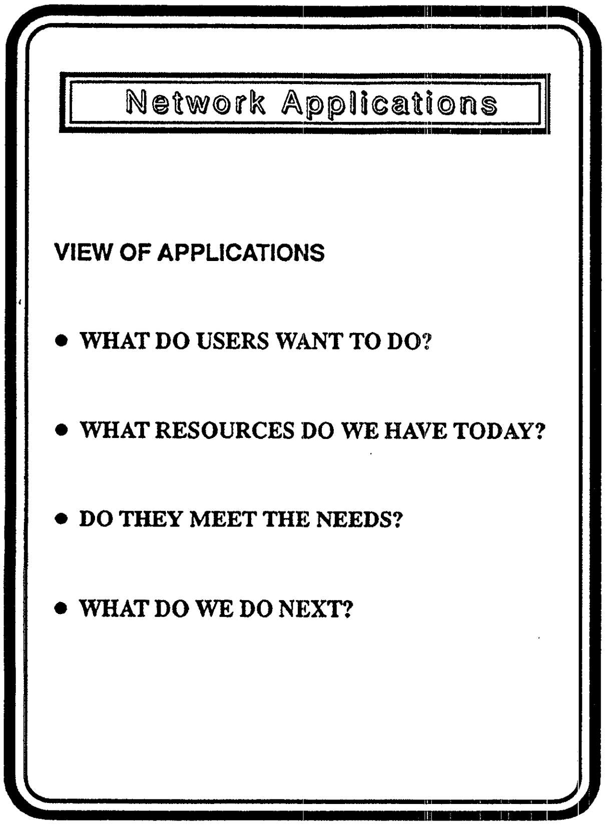

Network Applications

VIEW OF APPLICATIONS

- WHAT DO USERS WANT TO DO?
- WHAT RESOURCES DO WE HAVE TODAY?
- DO THEY MEET THE NEEDS?
- WHAT DO WE DO NEXT?

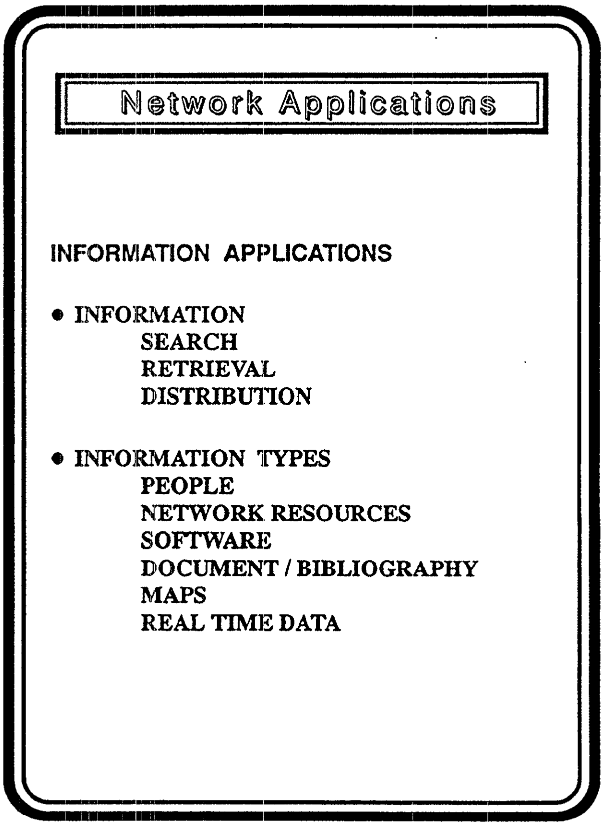

Network Applications

INFORMATION APPLICATIONS

- INFORMATION
  - SEARCH
  - RETRIEVAL
  - DISTRIBUTION
- INFORMATION TYPES
  - PEOPLE
  - NETWORK RESOURCES
  - SOFTWARE
  - DOCUMENT / BIBLIOGRAPHY
  - MAPS
  - REAL TIME DATA

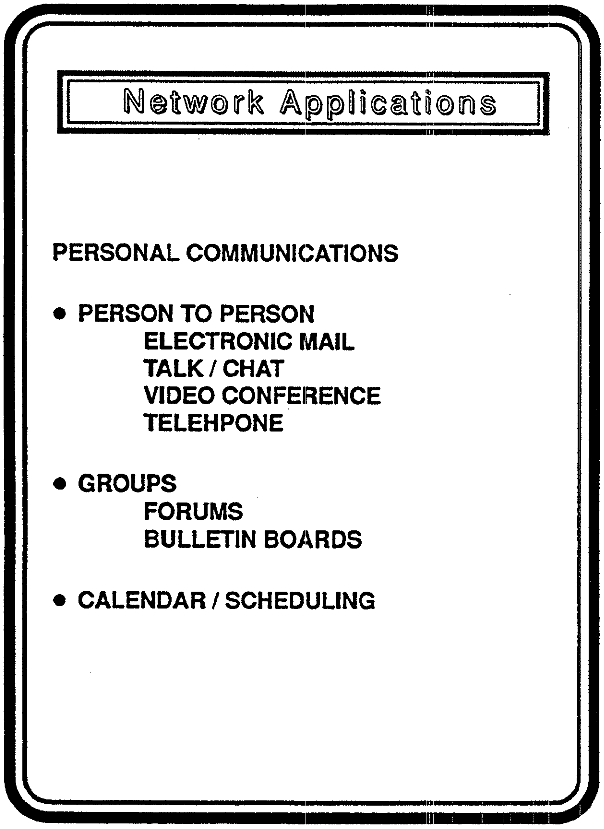

Network Applications

PERSONAL COMMUNICATIONS

- PERSON TO PERSON
  - ELECTRONIC MAIL
  - TALK / CHAT
  - VIDEO CONFERENCE
  - TELEHPONE
- GROUPS
  - FORUMS
  - BULLETIN BOARDS
- CALENDAR / SCHEDULING

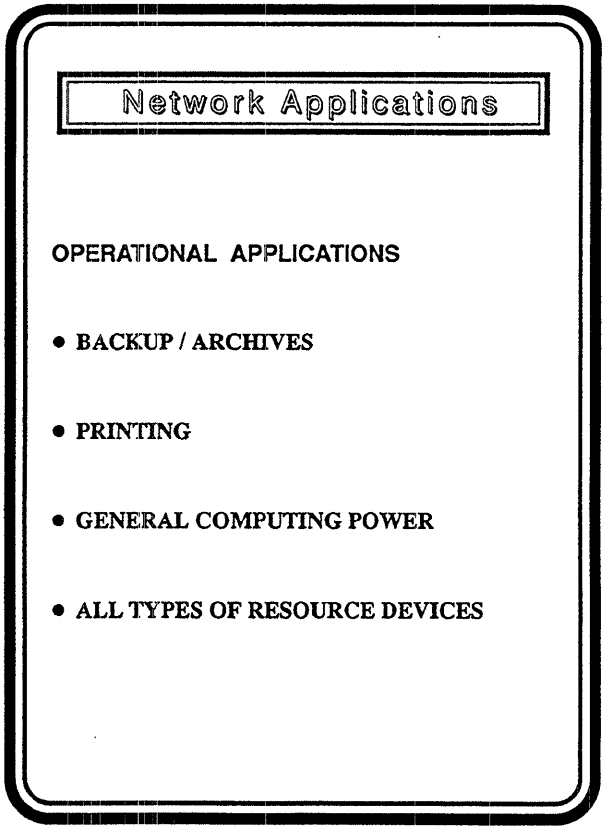

Network Applications

OPERATIONAL APPLICATIONS

- BACKUP / ARCHIVES
- PRINTING
- GENERAL COMPUTING POWER
- ALL TYPES OF RESOURCE DEVICES

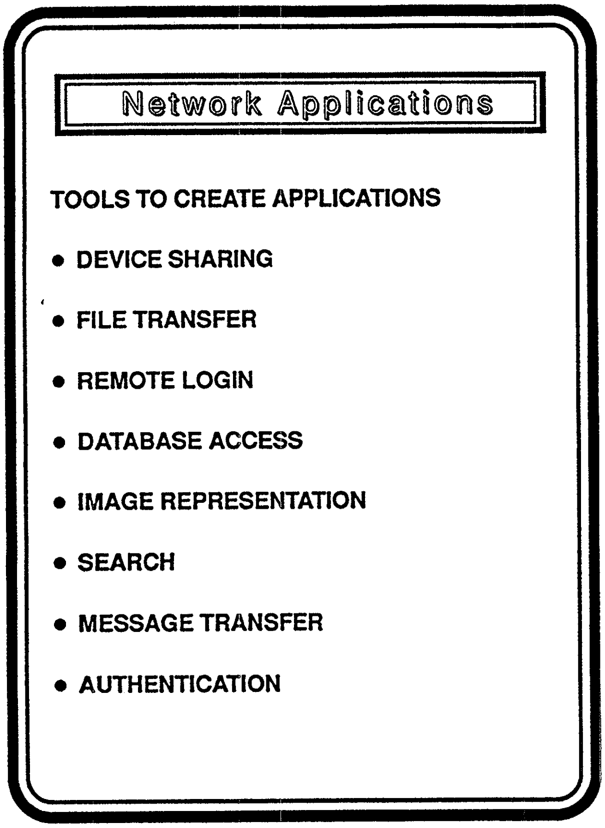

Network Applications

TOOLS TO CREATE APPLICATIONS

- DEVICE SHARING
- FILE TRANSFER
- REMOTE LOGIN
- DATABASE ACCESS
- IMAGE REPRESENTATION
- SEARCH
- MESSAGE TRANSFER
- AUTHENTICATION

3.1.1 Network FAX (netfax)

Charter

Chairperson:
Mark Needleman, mhn@stubbs.ucop.edu

Mailing Lists:
General Discussion: netfax@stubbs.ucop.edu
To Subscribe: netfax-request@stubbs.ucop.edu

Description of Working Group:

The Network Fax Working group is chartered to explore issues involved with the transmission and receipt of facsimile across TCP/IP networks and to develop recommended standards for facsimile transmission across the Internet. The group is also intended to serve as a coordinating forum for people doing experimentation in this area to attempt to maximise the possibity for interoperability among network fax projects.

Among the issues that need to be resolved are what actual protocol or protocols will be used to do the actual data transmission between hosts, architectural models for the integration of fax machines into the existing internet, what types of data encoding should be supported, how IP host address to phone number conversion should be done and associated issues of routing, and develeopment of a gateway system that will allow existing Group 3 and Group 4 fax machines to operate in a network enviornment.

It is expected that the output of the working group will be one or more RFC's documenting recommended solutions to the above questions and possibly also describing some actual implementations. The life of the working group is expected to be 18-24 months.

It is also hoped th at some fax vendors, as well as the networking community and fax gateway developers, will be brought into the effort.

Goals and Milestones:

Aug 1990	Review and approve charter making any changes deemed necessary. Refine definition of scope of work to be accomplished and intial set of RFC's to be developed. Begin working on framework for solution.

Mar 1991 Continue work on definition of issues and protocols. Work to be conducted on mailing list.

Aug 1991 First draft of RFC to be completed. To be discussed at IETF meeting and revised as necessary.

Dec 1991 Continue revisions based on comments received and if ok give to IESG for publication as RFC.

Mar 1992 Overlapping with activities listed above may be implementations based on ideas and work done by the working group. If so revise RFC to include knowledge gained from such implementations.

3.1.2 Network Printing Protocol (npp)

Charter

Chairperson:
Leo McLaughlin, `ljm@twg.com`

Mailing Lists:
General Discussion: `print-wg@pluto.dss.com`
To Subscribe: `print-wg-request@pluto.dss.com`

Description of Working Group:

The Network Printing Working Group has the goal of pursuing those issues which will facilitate the use of printers in an internetworking environment. In pursuit of this goal it is expected that we will present one or more printing protocols to be considered as standards in the Internet community.

This working group has a number of specific objectives. To provide a draft RFC which will describe the LPR protocol. To describe printing specific issues on topics currently under discussion within other working groups (e.g., security and dynamic host configuration), to present our concerns to those working groups, and to examine printing protocols which exist or are currently under development and assess their applicability to Internet-wide use, suggesting changes if necessary.

Goals and Milestones:

Done	Review and approve the charter, making any changes deemed necessary. Review the problems of printing in the Internet.
Apr 1990	Write draft LPR specification.
May 1990	Discuss and review the draft LPR specification. Discuss long-range printing issues in the Internet. Review status of Palladium print system at Project Athena.
May 1990	Submit final LPR specification including changes suggested at the May IETF. Discuss document on mailing list.
Jun 1990	Submit LPR specification as an RFC and standard.

Jul 1990 Write description of the Palladium printing protocol (2.0) in RFC format.

Aug 1990 Discuss and review the draft Palladium RFC.

CURRENT MEETING REPORT

Reported by Leo McLaughlin/ Wollongong

Minutes

Two primary tasks were accomplished at the May IETF.

One, the specification for LPR will be modified to support cacheless clients and servers by allowing the control file to be submitted before the data file and by allowing graceful end of connection instead of a field length to show end of file. In addition, control file lines beginning with 'A' and 'a' will be reserved for possible future use with Palladium.

Two, that use of Palladium, the Project Athena printing protocol, was seen as a good, long term, goal for printing in the Internet. As part of the Palladium 2.0 efforts currently under way, the Project Athena implementation (likely to be the future reference implementation) will be modified to support LPR clients. An RFC describing the printing protocol specific portions of Palladium is forthcoming.

Administrative Details

The mailing list of this working group is `print-wg@pluto.dss.com`, requests should be sent to `print-wg-request@pluto.dss.com`. We will be meeting in British Columbia.

ATTENDEES

Fred Bohle	fab@saturn.acc.com
Dave Borman	dab@cray.com
David Burdelski	daveb@ftp.com
Andrew Cherenson	arc@sgi.com
Bruce Crabill	bruce@umdd.umd.edu
Peter Dicamillo	cmsmaint@brownvm.brown.edu
Roger Fajman	raf@cu.nih.gov
Brian Handspicker	bd@vines.dec.com
Richard Hart	hart@decvax.dec.com
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Tom Holodnik	tjh@andrew.cmu.edu
Josh Littlefield	josh@cayman.com
John Loverso	loverso@xylogics.com
Matthew Nocifore	matthew@cupr.ocs.drexel.edu
Michael Petry	petry@trantor.umd.edu
Richard Smith	smiddy@dds.com
John Veizades	veizades@apple.com
Aileen Yuan	aileen@gateway.mitre.org

3.1.3 TELNET (telnet)

Charter

Chairperson:
Dave Borman, dab@cray.com

Mailing Lists:
General Discussion: telnet-ietf@cray.com
To Subscribe: telnet-ietf-request@cray.com

Description of Working Group:

The TELNET Working Group is to look at RFC 854, "Telnet Protocol Specification", in light of the last 6 years of technical advancements, and determine if it is still accurate with how the TELNET protocol is being used today. This group will also look at all the numerous TELNET options, and decide which of them are still germane to current day implementations of the TELNET protocol.

- Re-issue RFC 854 to reflect current knowledge and usage of the TELNET protocol.
- Create RFCs for new TELNET options to clarify or fill in any missing voids in the current option set. Specifically:
  - Environment variable passing
  - Authentication
  - Encryption
  - Compression
- Act as a clearing-house for all proposed RFCs that deal with the TELNET protocol.

Goals and Milestones:

Done	Write an environment option
Dec 1990	Write an authentication option
Dec 1990	Write an encryption option
Mar 1991	Rewrite RFC 854

CURRENT MEETING REPORT

Reported by David A. Borman/ Cray Research, Inc.

AGENDA

1. Linemode Option
2. Environment Option
3. Authentication Option
4. Encryption Option
5. Compression Option
6. TN3270 Option

The TN3270 option was not discussed. A discussion of TN3270 over Telnet was held over supper Wednesday evening by the interested parties, the minutes of that meeting are attached to the end of this report.

MINUTES

The COMPRESSION option was only briefly mentioned. When doing both encryption and compression, it is important that the sender apply the compression option before the encryption option, and that the receiver decrypt and then decompress. Both the ENCRYPTION and COMPRESSION documents will be modified to reflect this.

The LINEMODE option is currently a "proposed standard", RFC 1116. We discussed some additions to the option, two new mode bits and eight new special character definitions. After a brief explanation and minimal discussion, the two new mode bits (SOFT_TAB and LIT_ECHO) were accepted.

SOFT_TAB When set, the client side should expand the Horizontal Tab (HT) code, USASCII 9, into the appropriate number of spaces to move the printer to the next horizontal tab stop. When unset, the client side should allow the Hor-izontal Tab code to pass through un-modified.

LIT_ECHO When set, if the client side is echoing a non-printable character that the user has typed to the users screen, the character should be echoed as the literal character. If the LIT_ECHO bit is not set, then the client side may echo the character in any manner that it desires. (Many systems echo unprintable characters as two character sequences, for example, they will echo "^A" for an ASCII 1 value.)

Several new special characters, for systems that support in-line display editing of the command line, were proposed.

SLC_MCL Move cursor one character left. When visual editing is supported, this is the character which, when typed, will move the cursor one character to the left in the display.

SLC_MCR Move cursor one character right. When visual editing is supported, this is the character that, when typed, will move the cursor one character to the right in the display.

SLC_MCWL Move cursor one word left. When visual editing is supported, this is the character that, when typed, will move the cursor one word to the left in the display.

SLC_MCWR Move cursor one word right. When visual editing is supported, this is the character that, when typed, will move the cursor one word to the right in the display.

SLC_MCBOL Move cursor to the beginning of the line. When visual editing is supported, this is the character that, when typed, will move the cursor to the beginning of the line that is being edited.

SLC_MCEOL Move cursor to the end of the line. When visual editing is supported, this is the character that, when typed, will move the cursor to the end of the line that is being edited.

SLC_INSRT Toggel insert versus overstrike mode. When visual editing is supported, this is the character that, when typed, will toggle whether normal characters should be inserted into the display, or should overwrite characters the current display.

SLC_EWR Erase word to the right. When visual editing is sup- ported, this is the character that, when typed, will erase one word to the right of the cursor.

It was decided SLC_INSRT would be split into two values:

SLC_INSRT Enter character insert mode. When visual editing is supported, this is the character that, when typed, indicates that normal characters should be inserted into the display at the current cursor position.

SLC_OVER Enter character overstrike mode. When visual editing is supported, this is the character that, when typed, will indicate that normal characters should overwrite characters currently displayed.

If the SLC_INSRT and SLC_OVER values are set to the same value, than that value is to act as a toggle between insert and overstrike mode.

Three other special characters were added to round out the set:

SLC_ECR Erase one character to the right.

SLC_EBOL Erase from the current cursor position to the beginning of the line.

SLC_EEOL Erase from the current cursor position to the end of the line.

Also, in the current document, the SLC_EW description states what a "word" is:

> "... a word is defined to be (optionally) whitespace (tab or space characters), and a string of characters up to, but not including, whitespace or line delimiters."

With the addition of SLC_EWR, SLC_MCWL and SLC_MCWR, it was felt that this definition of "word" was no longer accurate. Rather than try to define what a "word" is, it was decided that we would remove this definition from the document, and put in some comments on why a "word" is not defined (to allow dissimilar systems to interoperate).

With these changes, it was recommended by the group that the LINEMODE option be re-issued as a "Draft Standard".

The ENVIRON option was discussed. A proposal was put forward to have the ENVIRON option issued as an RFC, as a "proposed standard". Section 6, "Well Known Variables" was discussed at length. People disagreed what the user account name variable should be, USER or USERNAME (some systems use LOGNAME). The group could not agree on what would be the best names for well known names, whether they should have a consistent format, (e.g., a common prefix) or whether there should be a common prefix for user-defined variables. Because resolution was not reached, it was decided that we would strike Section 6 from the document, but leave the variable names in the example section. We agreed that well known names could be added later if consensus was reached on the naming scheme.

Other changes: Explicitly state that the default set of variables is implementation dependent. Reword the motivation section to not be so "environment variabable" biased, since this option is used to pass arbitrary information, which happens to include environment variables. A "Security Considerations" section will be added, Jeff Schiller has agreed to write this.

The ENCRYPT option was briefly discussed. Comments that Steve Bellovin had made were touched upon. It was agreed that when encryption is being done, telnet options will be inserted BEFORE encryption is begun. We also need to add some comments about key management, and provide sub-options to allow for any initial negotiation required in a particular encryption scheme.

The rest of the meeting focused on the AUTHENTICATION option. There was some major re-structuring of how the option works. Previously, DO/WILL AUTHENTICATION was sent in each direction for each direction that authentication was desired. Unfortunately, this breaks down if the authentication scheme has a third method; mutual authentication. It was decided that enabling the AUTHENTICATION option in either direction enables authentication. A definition of "server" and "client" will be added ("server" is the side that did the "passive" TCP open, and client is the side that did the "active" TCP open).

The "server" sends the "IAC SB AUTHENTICATION SEND ... IAC SE" command, and the client sends the "... IS ..." command. The server my optionally respond to the IS with a REPLY, and the client may optionally respond to a REPLY with another IS. This way, the client and server may do as many exchanges of information as necessary for the particular authentication scheme being used.

The "authentication-type" sent in SEND, IS and REPLY commands is now a triplet, <type><authenticator/authenticatee><one-way/mutual>. Several things needed to be determined: i.e., who will initiate the authentication, who is being authenticated, and in which direction (server authenticates client, client authenticates server, client and server authenticate each other). We decided that the server side always initiates the authentication procedure (only the server can send a SEND command). The other two parts indicate how the authentication is being done. Authenticator/authenticatee indicates whether the server is authenticating the client, or the client is authenticating the server. One-way/mutual is whether the authentication is only happening on one side, or whether both sides authenticate each other. (Authenticator-mutual and authenticatee-mutual allow the authentication scheme to distinguish who initiates authentication.)

The list of authentication-types sent with the SEND command MUST be an ordered list of preferences of the server, so that the client can reliably know which authentication scheme is preferred.

There was also some discussion about what to do with normal data that comes across the telnet data stream before the authentication is completed. What happens to the data will be implementation dependent. Telnet options received during authentication must be processed in the normal manner, but an implementation might choose to refuse or delay the effect of certain options until the authentication has been completed.

It was also decided to add a generic LOGIN authentication type, which is the normal login:/password: prompting.

A security consideration section will be added. It will state that successfully authentication does not imply that the entire session is secure; the connection might still be taken over after the authentication is done.

There is a reference to the "Assigned Numbers" RFC that will be removed.

For action items, Dave Borman will integrate these changes into the Option drafts, and send them off to the internet-drafts directory; Russ Hobby will be notified when the LINEMODE and ENVIRON options are ready, so that they can be pushed on to being issued as RFCs.

Minutes of Dinner Meeting

Minutes of the special interest group/dinner that met at the Holiday Inn at 7:00 PM on 5/2/90.

The group discussed the current mechanism for specifying the use of and problems with 3270 data-streams within a TELNET session. After some discussion, it was decided to write a new RFC for specifying 3270 mode. Features of this RFC would include:

- Single option for negotiating 3270 mode.
- Information about terminal characteristics (size, color support, etc.) defined within the 3270 Data-Stream using the Write Structured Field Query Reply facility negates the need for the use of the TERMINAL-TYPE option.
- New option implies TRANSMIT-BINARY, which does not need to be separately negotiated.

- Uses a TLV type structure to encapsulate the 3270 Data-Stream. This allows optional items to be sent and received. Examples of this include an indicator that the following data has a READ command chained to it, and for for 3270 type printers to be able to send their completion status back to the server. This mechanism also allows for future extensions.
- Within a given TLV (Type, Length, Value) structure, the data is not IAC stuffed. TELNET commands and options may occur between individual TLV structures.
- The new option is negotiated only by the server. Since 3270 Data Streams require both directions to be in the mode, it didn't seem necessary to require it to be negotiated in both directions. This will simplify server and client implementations.
- Allow the 3270 Data Stream to be unnegotiated and renegotiated as needed by the server.
- Require clients to support SNA and non-SNA commands.
- No longer requires the EOR option or the use of the EOF TELNET command.
- Spent significant time discussing printing issues. Decided to write a seperate RFC on this issue since there appear to be several ideas on how this could be solved.

ATTENDEES

Fred Bohle	fab@saturn.acc.com
David Borman	dab@cray.com
Bruce Crabill	bruce@umdd.umd.edu
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Roger Fajman	raf@cu.nih.gov
James Galvin	galvin@tis.com
Mike Horowitz	mah@shiva.com
Phil Karn	Karn@Thumper.Bellcore.Com
John LoVerso	loverso@xylogics.com
Louis Mamakos	louie@trantor.umd.edu
Greg Minshall	minshall@kinetics.kinetics.com
Gerard Newman	gkn@sds.sdsc.edu
Michael Petry	petry@trantor.umd.edu
Jeffrey Schiller	jis@athena.mit.edu
Frank Solensky	solensky@interlan.interlan.com
Ted Soo-Hoo	soo-hoo@dg_rtp.dg.com
Peter Vinsel	farcomp!pcv@apple.com

Participants of the dinner meeting were:

Fred Bohle	fab@saturn.acc.com
David Borman	dab@cray.com
Bruce Crabill	bruce@umdd.umd.edu
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Roger Fajman	raf@cu.nih.gov
Yakov Rekhter	yackov@ibm.com

3.2 Host and User Services Area

Director: Craig Partridge/BBN

Host Services

Three WG's in the host area met: User Connectivity Problems, Dynamic Host Configuration and Distributed File Systems.

Because Dan Long couldn't make the meeting, Kent England chaired the User Connectivity Problems WG meeting (thanks to Kent for helping out!). The WG considered slightly different models for user connectivity proposed by Elise Gerich, Karen Bowers and Craig Partridge. The group discussed various issues raised by the proposals. One key decision was that a coordinated trouble ticket system seems essential to all three schemes and Matt Mathis volunteered to write up some discussion of the issues (which he has done). Another point was that we need to understand the "boundary" of the system – i.e., who is inside (and responsible for fixing things) and who is outside (needing repairs to be made).

Dynamic Host Configuration made good progress on trying to come to closure on key issues so that an RFC can be issued this year. In particular, the WG decided on the parameters necessary to configure the client's network layer and decided to base its protocol on BOOTP. The WG is currently looking at address assignment mechanisms in servers.

The Distributed File Systems WG spent the meeting devoted to a lengthy examination of the NFS protocols, and generated a variety of recommendations and issues related to running NFS over TCP/IP.

User Services:

Reported by Joyce Reynolds

User-Doc WG - Coming to a close

Chaired by Tracy LaQuey and Karen Roubicek.

The User-Doc Bibliography is ready for the Internet Draft Process. Final changes or amendments to the Bibliography have a deadline date of May 15th.

After the Internet-Draft process, to the RFC publication, the User-Doc WG will terminate, and go back into the USWG.

We are pleased that in just a 12 month time period the User-Doc WG produced their

NISI - Reinstated

Chaired by Dana Sitzler

This IETF is NISI's (Network Information Services Infrastructure) first meeting since its reinstatement after the last IETF at Tallahassee.

NISI focused on discussion of the old NISI charter and a survey of "where we are now" (i.e., a survey of existing informational types, retrieval mechanisms and current NIC specialties and relationships), specifically, how to get information to people. A draft will be sent to the NISI mailing list.

SSPHWG - Security Area/User Services Area combined efforts

Chaired by J. Paul Holbrook and Joyce K. Reynolds

The SSPHWG (Site Security Policy Handbook WG) held its first meeting in Pittsburgh. It had a great turnout of thirty people, with a good mixture of USWG members and Security Area members. Primary meeting time focused on development of an outline for a Handbook. Twenty-two bullets were developed in a scratch outline. Volunteers will take on the task of developing a draft outline to be presented at the next SSPHWG meeting. We have a very ambitious schedule, as this WG would like to have a completed Handbook available for distribution by the end of this year.

This WG is the first to combine the efforts of two separate IETF Areas. The response to this had been successful. Steve Crocker thinks it's a "neat" idea. It IS okay to "cross the streams" between the IETF Areas. Other Areas and WGs are encouraged to follow suit, if they feel the need.

The next meeting of the SSPHWG will be held at USC/Information Sciences Institute in Marina del Rey on Tuesday, June 12th.

USWG - Running at its peak

Chaired by Joyce K. Reynolds

USWG Announcements:

- FYI RFC sub-series start-up
- NOCTOOLS publication (RFC1147, FYI2)
- NOCTOOLS was historically an offspring of the USWG.

Agenda items included:

- Distribution and Announcement Handbook
- Question and Answer Mailing List
- Intro Packages
  - what exists
  - what is needed

3.2.1 Distributed File Systems (dfs)

Charter

Chairperson:

Peter Honeyman, honey@citi.umich.edu

Mailing Lists:

General Discussion: dfs-wg@citi.umich.edu
To Subscribe: dfs-wg-request@citi.umich.edu

Description of Working Group:

Trans- and inter-continental distributed file systems are upon us. The consequences to the Internet of distributed file system protocol design and implementation decisions are sufficiently dire that we need to investigate whether the protocols being deployed are really suitable for use on the Internet. There's some evidence that the opposite is true, e.g., some DFS protocols don't checksum their data, don't use reasonable MTUs, don't offer credible authentication or authorization services, don't attempt to avoid congestion, etc. Accordingly, a working group on DFS has been formed by the IETF. The WG will attempt to define guidelines for ways that distributed file systems should make use of the network, and to consider whether any existing distributed file systems are appropriate candidates for Internet standardization. The WG will also take a look at the various file system protocols to see whether they make data more vulnerable. This is a problem that is especially severe for Internet users, and a place where the IETF may wish to exert some influence, both on vendor offerings and user expectations.

Goals and Milestones:

May 1990	generate an RFC with guidelines that define appropriate behavior of distributed file systems in an internet environment.

CURRENT MEETING REPORT

Reported by Peter Honeyman/ University of Michigan

Minutes

At this meeting, attention was focused on NFS. The consensus was that it will be most useful in the near term to draft a "survival guide" for NFS. The audience for this guide will be vendors and system administrators.

Suggested recommendations were discussed. Items that can be addressed only by implementors are noted. Some items suggest coordination with the NOC tools.

- Avoid packet retransmission
  - Soft mounts vs. hard mounts
  - Adjust timeout parameters to meet local conditions
  - Transaction ID caching (implementation)
  - Adaptive retransmission strategy (implementation)
- Avoid IP fragmentation : Adjust read and write sizes to meet local conditions
- Ensure reliable data transfer: Use UDP checksum for long-haul
- Privacy issues
  - Reserved socket myth
  - Mutual distrust among client and server (implementation)
  - Periodic FSIRAND (NOC tools)
  - Setuid handling
  - IP address verification at mount time
  - IP address verification at access time (implementation)
  - Root and anonymous mapping
  - Generalized mapping (implementation)
- System management
  - NFSSTAT and NFSWATCH (NOC tools)
  - SNMP for NFS (implementation)
  - Export controls
  - Cache timeout management

Since Ethernet checksum can obviate UDP checksum, a suggestion was made that UDP checksum be a mount option. This may not be practical, since most NFS servers are running on an operating system for which UDP checksum is either always enabled or always disabled. The consensus seems to be that correctness is more important than performance. i.e., UDP checksum should always be enabled. It was reported that in some vendors' operating systems, it is impossible to turn on UDP checksum.

There was further discussion of the protocol for Kerberos integration with NFS.

ATTENDEES

David Burdelski	daveb@ftp.com
Andrew Cherenson	arc@sgi.com
Sailesh Chutani	chutani@transarc.com
Bruce Crabill	bruce@umdd.umd.edu
Peter DiCamillo	smsmaint@brownvm.brown.edu
Craig Everhart	cfe@transarc.com
Dennis Ferguson	dennis@gw.ccie.utoronto.ca
Fred Glover	fglover@decvax.dec.com
Olafur Gudmundson	ocud@cs.umd.edu
Peter Honeyman	honey@citi.umich.edu
Steve Hubert	hubert@cac.washington.edu
Tim Hunter	thunter@allegum
Steven Hunter	hunter@ccc.mfecc.arpa
Josh Littlefield	jost@cayman.com
Tony Mason	mason@transarc.com
Leo McLaughlin	ljm@twg.com
Greg Minshall	minshall@kinetics.kinetics.com
Jeffrey Mogul	mogul@decwrl.dec.com
Dan Nydick	nydick@psc.edu
Brad Parker	brad@cayman.com
Drew Perkins	ddp@andrew.cmu.edu
Joel Replogle	replogle@ncsa.uiuc.edu
Bob Sidebotham	bobo@andrew.cmu.edu
Ted Soo-Hoo	soo-hoo@dj-rtp.dg.com
Brad Strand	bstrand@cray.com

3.2.2 Dynamic Host Configuration (dhc)

Charter

Chairperson:
Ralph Droms, droms@sol.bucknell.edu

Mailing Lists:
General Discussion: host-conf@sol.bucknell.edu
To Subscribe: host-conf-request@sol.bucknell.edu

Description of Working Group:

The purpose of this working group is the investigation of network configuration and reconfiguration management. We will determine those configuration functions that can be automated, such as Internet address assignment, gateway discovery and resource location, and those which cannot ne automated (i.e., those that must be managed by network administrators).

Goals and Milestones:

Jun 1990	We will identify (in the spirit of the Gateway Requirements and Host Requirements RFCs) the information required for hosts and gateways to: Exchange Internet packets with other hosts, Obtain packet routing information, Access the Domain Name System, and Access other local and remote services.
Jun 1990	We will summarize those mechanisms already in place for managing the information identified by objective 1.
Jan 1991	We will suggest new mechanisms to manage the information identified by objective 1.
Jan 1991	Having established what information and mechanisms are required for host operation, we will examine specific scenarios of dynamic host configuration and reconfiguration, and show how those scenarios can be resolved using existing or proposed management mechanisms.

CURRENT MEETING REPORT

Reported by Ralph Droms/ Bucknell

This meeting of the DHC WG concentrated on the details of the proposed DHC protocol. Specifically, the WG concentrated on the DHC protocol as used to initially configure the client's network layer. The WG agreed that the following parameters should be configured:

- IP address
- Subnet mask
- Broadcast address
- (Non-default or unusual) MTU - which may be required by some kinds of network hardware

The WG has further agreed to base the DHC protocol on the BOOTP protocol, as extended by R. L. Morgan. The agenda items for this meeting, then, included the definition of the following:

- Client behavior within the protocol
- Server behavior
- Router or other forwarding agent behavior
- Protocol message formats

There are two primary problems to be solved by the client: first, the client must decide which of possibly several sources of configuration information to use and second, the client must decide which IP address to use if given a range of addresses to choose from. The client may get configuration information from a local cache or from a DHC protocol server. If no configuration information is available (the "genesis state"), the client should use a default configuration that allows interoperation with other clients on the same local net.

Greg Minshall presented an algorithm (included with this report) that was discussed at the meeting. The genesis state was discussed at some length. The WG agreed that a client in the genesis state should use a distinguished network number, defined so that routers will never forward packets with the distinguished network number. This distinguished network number will allow interoperation between hosts on an isolated network, with no danger of genesis state packets leaking onto the internet if the isolated network becomes attached to an internet at some later time. The WG also discussed problems with the transition from the genesis state to a normally configured state. If an isolated net becomes attached while hosts are in genesis state, the hosts will either have to restart to obtain correct configuration parameters, or must be able to support interoperation with two logical nets on the same interface (both the genesis state network and the "real" network).

The WG briefly discussed router behavior. We need to find out from router vendors about the details of existing BOOTP implementations so the WG can assess the impact of changes to the BOOTP protocol on existing implementations and determine if a formal description of BOOTP forwarding agent behavior needs to be written.

The final point of discussion sparked some real controversy. Before this meeting, the WG had discussed the IP assignment mechanism as an extension of the MIT and Morgan/BOOTP mechanisms, in which a client is provided a range of IP addresses from which it can choose a preferred IP address. At this meeting, an alternative proposal was presented, in which BOOTP servers were presumed to have sufficient knowledge of the network configuration so as to be able to determine and allocate a single IP address to a client. The presumption was that such a dynamic allocation mechanism would make the client code much simpler (in fact, existing BOOTP client code would work unchanged) at an acceptable cost in server complexity. The dissenting opinion was that the increased server complexity was not worth the simplification in the client code.

As neither side had anything in writing, the WG had some difficulty in arguing the relative merits of the two mechanisms. The WG chair has scheduled a meeting for June 8 in which several of the participants in the WG discussion will present written descriptions of the two mechanisms for discussion.

ATTENDEES

Douglas Bagnall	bagnall_d@apollo.hp.com
Terry Braun	tab@kinetics.com
Andrew Cherenson	arc@sgi.com
Peter DiCamillo	cmsmainto@brownvm.brown.edu
Hunaid Engineer	hunaid@opus.cray.com
Roger Fajman	raf@cu.nih.gov
Metin Feridun	mferidun@bbn.com
Karen Frisa	karen@kinetics.com
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Tom Holodnik	tjh@andrew.cmu.edu
Mike Horowitz	mah@shiva.com
Leo McLaughter	ljm@twg.com
Greg Minshall	minshall@kinetics.kinetics.com
Jeffrey Mogul	mogul@decwrl.dec.com
Michael Reilly	reilly@nsl.dec.com
Jeffrey Schiller	jis@athena.mit.edu
Tim Seaver	tas@mcnc.org
Ted Soo-Hoo	soo-hoo@dg-vtp.dg.com
John Veizades	veizades@apple.com
Steve Waldbusser	sw01@andrew.cmu.edu
Jonathan Wenocur	jhw@shiva.com

CURRENT MEETING REPORT

Reported by Ralph Droms/ Bucknell University

Introduction

John Veizades, Jeff Mogul, Greg Minshall, Bob Morgan, Leo McLaughlin and Ralph Droms attended a "mid-term" meeting of the Dynamic Host Configuration working group. Jeff Mogul was kind enough to host the meeting at DEC's Western Research Lab. The purpose of the meeting was to discuss the mechanism for network address allocation proposed at the Pittsburgh IETF plenary. The participants were chosen to represent the two mechanisms as presented at the Pittsburgh meeting. The time and location were chosen for the convenience of the participants.

The allocation mechanism under discussion at this meeting describes the way in which DHC servers allocate and transmit network addresses to DHC clients. This new mechanism was first proposed by Leo and Jeff at the Pittsburgh meeting. In this mechanism, the DHC servers allocate and return a single network address to a requesting client.

Discussion

The DHC WG generally agrees that the new DHC protocol should be based on the existing BOOTP protocol. The primary motivations behind this decision are the desire to capture BOOTP forwarding agent code in existing routers and the desire to avoid inventing a new protocol when an existing protocol can be used.

The WG further agrees that the new protocol should be first defined to carry network layer configuration parameters to the client: network address, subnet mask, broadcast address and local network MTU. The question arises: how shall the DHC server select a network address to return to the client? There are several points on which one can compare the two network address allocation mechanisms discussed in the introduction:

- Relative complexity of client and server code
- Accuracy/correctness of allocated addresses
- Compatibility with existing BOOTP clients
- Ability to maintain coherent distributed information about allocated addresses

The explicit allocation mechanism has appeal because it captures existing BOOTP clients and because it can make the client code much simpler. However, at the Pittsburgh meeting there was much discussion about whether the central allocation mechanism could be made to work; would it be too complex and would it be possible to maintain the global database required for distributed allocation of addresses?

New Mechanism

At the June meeting, we developed an outline of the specific address mechanism, which we present for comment here. From the client's point of view, the new DHC protocol works the same as the existing BOOTP mechanism. In fact, we expect existing BOOTP clients to interoperate with DHC servers without requiring any change to initialization software. There are some new, optional transactions that optimize the interaction between DHC clients and servers:

- Client sends DHC request for network parameters
- Server sends response with network parameters and explicit network address (Note: This assumes the client receives at least one reply. If no replies arrive, the client may, at the discretion of local administration, enter "genesis state" [see below for details].)
- (Opt.) Client updates local network ARP caches with an ARP broadcast reply
- (Opt.) Client releases unused addresses from duplicate server responses
- (Opt.) Client releases selected address during orderly close

Problem Areas with Explicit Allocation

The first problem area encountered by a server when explicitly allocating a specific network address rather than a range of addresses is determining which addresses are already in use and which may be allocated or reallocated. Because the server may not be on the same subnet as the client, the server must use an ICMP echo message to probe for hosts already using a specific address. Thus all participants using network addresses in the dynamic allocation range (whether statically or dynamically allocated) must implement ICMP echo message processing.

Some hosts, while implementing ICMP echo processing, may go into a state where ICMP echo requests are ignored for extended periods. The client request protocol includes two new extensions to help the server handle such clients:

- "brain damage" - indicating that the client may ignore ICMP requests
- "reserve forever" - requesting permanent allocation of a network address

To meet the goal of reissuing the same network address to a host whenever possible, while allowing more hosts on a subnet than addresses available for allocation (obviously, not all hosts can be active simultaneously), the server must be able to timeout the allocation of an address to a host, and reuse addresses in LRU order. The optional "release address" message from the client to the server also helps the server determine when a network address may be reallocated

The second problem area, which was discussed at some length in Pittsburgh, is the mechanism through which multiple servers can coordinate the allocation of addresses.

We believe this is not a difficult problem. First, note that the problem only arises when multiple servers share responsibility for allocation of addresses on a single subnet. Second, the servers need not interact dynamically, after every network address is allocated. Rather, the address space for the target subnet can be partitioned among the servers, which each only allocate addresses from its own partition. Periodically, the servers exchange allocation information, possibly repartitioning the currently unallocated addresses to reflect client request load.

Genesis State

In the absence of any servers, clients may choose to enter "genesis state". This state is intended for use in small networks in which resources for support of DHC servers may not be available (the "dentist's office" network). In genesis state, the client picks an IP address, probes for any current use of that address and then defends the selected address using ARP. The genesis state mechanism looks much like the Athena NIP address allocation mechanism.

Conclusion

The problem areas in a protocol where network addresses are explicitly selected by a possibly remote server seem to be identifiable and can be surmounted by careful design of the protocol and server behavior. The advantages of explicit network address allocation appear to outweigh the disadvantages, and I recommend the DHC WG further investigate the new address allocation mechanism.

3.2.3 Internet User Population (iup)

Charter

Chairperson:
Craig Partridge, craig@nnsc.nsf.net

Mailing Lists:
General Discussion: ietf@venera.isi.edu
To Subscribe: ietf-request@venera.isi.edu

Description of Working Group:

To devise and carry out an experiment to estimate the size of the Internet user population.

Goals and Milestones:

Sep 1990	Write a description of the experimental procedure.
Jan 1991	Write an RFC that gives the results of the experiment.
TBD	Prepare an article for publication in a networking magazine.

3.2.4 Network Information Services Infrastructure (nisi)

Charter

Chairperson:
Dana Sitzler, dds@merit.edu

Mailing Lists:
General Discussion: nisi@merit.edu
To Subscribe: nisi-request@merit.edu

Description of Working Group:

The NISI WG will explore the requirements for common, shared Internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services "infrastructure" for the Internet. This effort will be a sub-group of the User Services WG and will coordinate closely with other efforts in the networking community.

Goals and Milestones:

Done First IETF meeting; review and approve charter. Begin information gathering process to write a short white paper to serve as a starting point for discussions on an Internet-wide information services infrastructure. This paper will document current available information and existing information retrieval tools.

Aug 1990 Review draft for phase 1 and begin discussions for completing the second phase which is to define a basic set of 'cooperative agreements' which will allow NICs to work together more effectively to serve users.

Jul 1990 Complete draft for phase 2 suggesting cooperative agreements for NICs.

CURRENT MEETING REPORT

Reported by Dana Sitzler/ Merit

Minutes:

The meeting began with some general guidelines as a framework for thinking about creating a network information services infrastructure. These 'guidelines' are listed below:

- Think future: Try not to limit our thinking to the current operations; think toward the NREN
- Information Services=User Services: Terminology doesn't matter; how are we going to provide the information and services that network users need
- Not just backbone NICs: Our discussions must account for any center providing support to users. This includes organizational NICs (campus)
- Users: Novice to Expert: We must accommodate a large variance in user knowledge, experience and comfort with networking
- Be Aware: Many on-going activities are related to this effort such as
  - User Connectivity WG
  - Distributed File Systems
  - NREN activity
  - Directory Services
  - Database Services

The meeting proceeded with a review of the draft charter. The objectives of the charter basically address three phases: Where are we; How do we work together to help users; How do we do it electronically. The charter of the working group is provided below:

Once the charter was approved, the group started brainstorming information to complete the first objective. The first objective deals with the current state of information services and will serve as a base on which to build. The topics of discussion were:

- what information is available now
- how is information accessed
- what information formats are available
- what kinds of services are NICs offering
- who are NICs serving

The group then discussed other projects on-going in the internet community about which more information is needed. These projects include directory services projects, library activity, work in distributed file systems, etc. Assignments were made for group members to gather information on these various projects.

Action Items:

Draft 'where are we' paper	Sitzler
Send FARNET info gathering survey to group	Roubicek
Investigate NYSERNet Directory Services Pilot	Sturtevant
Investigate Andrew Project	Moore
Investigate NFS	Sturtevant
Investigate Mercury Project	Roubicek
Investigate PSI Z39.50 Pilot	Hallgren
Investigate ISODE	Sturtevant
Investigate Database systems (info servers)	Stahl
Investigate Remote access DB systems	Carpenter

People investigating other projects will attempt to get an overview of how the system works, how it's used, what material it will handle, mechanisms for access, what the user interface is like, and how 'exploitable' it is.

ATTENDEES

Glee Cady	ghc@merit.edu
Jeffrey Carpenter	jjc@unix.cis.pitt.edu
Martyne Hallgren	martyne@tcgould.tn.cornell.edu
Ole Jacobsen	ole@csli.stanford.edu
Tracy Laquey	tracy@emx.utexas.edu
Marilyn Martin	martin@cdnnet.ca
Berlin Moore	prepnet@andrew.cmu.edu
Marc-Andre Pepin	pepin@crim.ca
Joyce Reynolds	jkrey@venera.isi.edu
Karen Roubicek	roubicek@nnsc.nsf.net
Pat Smith	psmith@merit.edu
Mary Stahl	stahl@nisc.sri.com
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
John Wobus	jmwobus@suvm.acs.syr.edu
Aileen Yuan	aileen@gateway.mitre.org

3.2.5 Special Host Requirements (shr)

Charter

Chairperson:
Bob Stewart, rlstewart@eng.xyplex.com

Mailing Lists:
General Discussion: ietf-hosts@nnsc.nsf.net
To Subscribe: ietf-hosts-request@nnsc.nsf.net

Description of Working Group:

The Special-purpose Host Requirements working group is chartered to clarify application of the Host Requirements RFCs (1122 and 1123) to systems that are technically hosts but are not intended to support general network applications. These special-purpose hosts include, for example, terminal servers (a "Telnet host"), or file servers (an "FTP host" or an "NFS host").

The Host Requirements RFCs address the typical, general-purpose system with a variety of applications and an open development environment, and give only passing consideration to special-purpose hosts. As a result, suppliers of special-purpose hosts must bend the truth or make excuses when users evaluate their products against the Requirements RFCs. Users must then decide whether such a product is in fact deficient or the requirements truely do not apply. This process creates work and confusion, and undermines the value of the RFCs. The commercial success of the Internet protocols and their use in increasingly unsophisticated environments exacerbates the problem.

The working group must define principles and examples for proper functional subsets of the general-purpose host and specifically state how such subsets affect the requirements. The working group must determine the balance between an exhaustive list of specific special-purpose hosts and philosphy that remains subject to debate. For the most part, it should be possible to base decisions on existing experience and implementations. The special-purpose requirements will be stated as differences from the existing RFCs, not replacements, and will refer rather than stand alone.

Since they define strict subsets of the Host Requirements RFCs, the Special-purpose Host Requirements appear to be an easier job and can be developed and stabilized within 8-12 months. Most of the group's business can be conducted over the Internet through email.

Goals and Milestones:

Jun 1990 Mailing list discussion of charter and collection of concerns.

Aug 1990 First IETF Meeting: discussion and final approval of charter; discussion and agreement on approach, including models, format, level and type of detail. Make writing assignments.

Oct 1990 First draft document.

Nov 1990 Second IETF Meeting: review first draft document, determine necessary revisions. Follow up discussion on mailing list.

Jan 1990 Revised document.

Feb 1990 Third IETF Meeting: make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail.

Apr 1991 Final draft document.

May 1991 Fourth IETF meeting: review final draft and if OK, give to IESG for publication as RFC.

3.2.6 User Connectivity (ucp)

Charter

Chairperson:
Dan Long, long@bbn.com

Mailing Lists:
General Discussion: ucp@nic.near.net
To Subscribe: ucp-request@nic.near.net

Description of Working Group:

The User Connectivity working group will study the problem of how to solve network users' end-to-end connectivity problems.

Goals and Milestones:

TBD	Define the issues that must be considered in establishing a reliable service to users of the Internet who are experiencing connectivity problems.
TBD	Write a document, addressing the above issues, which describes a workable mechanism for solving User Connectivity Problems. Address the above issues. Submit this document into the RFC pipeline as appropriate.

CURRENT MEETING REPORT

Reported by Ken England/ Boston University and Karen Roubicek/ BBN

Connectivity Tool Demonstrations

Metin Feridun made a brief announcement of demonstrations of the "Connectivity Tool" that he has been working on. The CT is designed to present a network detective of modest skills with a suite of analysis tools and built-in technique to simplify the process of tracking down internet connectivity.

Last Meeting

At the last (first) meeting of UCP-WG, Craig Partridge, Elise Gerich, and Karen Bowers each made presentations of a point of view on modeling the operations of the Internet. Unfortunately, none of these worthy thinkers was able to attend the IETF this time, so the host had to make due with unworthy re-presentations of these ideas and copious reference to notes from postings that these thinkers had made to the UCP list, prior to this meeting. Perhaps the original ideas came across anyway.

Craig Partridge's Model

Craig Partridge's model was reviewed. Karen Roubicek coined the term "UCP Central" to denote the national "center" with an 800 number, and this term was extended to include the following four elements of an architecture:

- UCP Central (the 800 number service)
- Site Entity
- A User (of this system under study)
- A Regional Entity (tentatively put forth for study)

Elise Gerich's Model

Elise identified some structure within the "UCP Central Entity" [note that terminology is deliberately vague, in order to avoid excessive connotative baggage -kwe]

In addition to recognizing Site and User Entities, like Craig's model, Elise put some structure to the UCP Central Entity, by postulating:

- National Center (we called it UCP Central)
- (Six) Regional Centers

and corresponding structure.

Karen Bowers' Model

Unfortunately for us, Karen has left the Internet community and was unable to write up a description of her model. The host was inadequate to the task of recalling her model, but members of the audience who had been impressed by her words last time recalled that Karen had allowed a richer connectivity from Site to Site or from Regional to Regional in her model.

Synthesis

Some common points arise from these models and beg some questions:

We must define a User Entity and consider how these Users, who may be end-users or may be lower level representatives of end-users, such as campus NOCniks, enter this system, how they interact with this system we are defining, and how their problems are staged and addressed. Assumptions of available tools and skills depends on who we assume the User to be.

We have to consider centralized (UCP Central) versus decentralized (Site/Regional Entity) issues, and clearly delineate responsibilities and interactions. We must consider the authority of the UCP Central and how it is derived.

We must consider the nature of the Site and Regional Entities; are they Network Operations Centers, or Network Information Centers, or both, or neither? Let us call these entities Network Service Centers (NSCs) for the moment, and withhold evaluation of what they really are.

General Discussion

Who is it that owns these facilities? Who are the players; the campuses, the regionals, the backbones, the commercial service providers, etc?

How will these entities; these Users and NSCs; be coordinated?

How do we resolve problems that the participants in this model cannot solve, such as host interoperability problems? Are there others that must get involved to solve these sorts of problems?

We need a means of filtering out chronic problems, ones that have been identified, but are not yet solved, or are unsolvable by our system.

Trouble Ticket Systems

Trouble ticket systems came up as something that seems to be an integral part of the solution of UCPs.

Matt Mathis commented that we need a protocol for managing ownership of trouble tickets, that we need some centralization for dealing with problems (UCP Central), but we must have filters so that UCP Central does not have to deal with too many routine problems. We also need to make sure that tickets don't "evaporate" and we could use a meta-UCP protocol for evaluating how well individual UCPs were handled by the system. We also need to discriminate equipment failures from infrastructure or engineering problems, which this system may not be able to handle. We also have to consider how the User is notified of progress on his Ticket.

Further Synthesis

What can we glean from what everyone has said so far?

1. We need to put a boundary around the problem; around the system we are trying to define.
2. "Users" are outside this system boundary. "Network Service Centers" are entities that are within the boundary of our system and our model.
3. Users need a "protocol" or procedure for how they interact with this system. Let us call this the P1 protocol; User-to-NSC.
4. NSCs need a "protocol" or procedure for how they interact among themselves. Let us call this the P2 protocol; NSC-to-NSC.
5. At a minimum, we need to define a "User", an "NSC", and the P1 and P2 protocols. Work in this direction will undoubtedly lead to further modeling requirements.

We need to consider at least these steps in the process:

- diagnosis of the problem
- the resolution process
- closure
- connectivity versus interoperability problems

Someone described the AT&T trouble ticket model, and noted that the person in the system that was "closest" to the end-user was responsible for updating the user on progress and for closure, but that the ticket database was centralized and centrally managed.

There was discussion of the P2 protocol and how it related to lines of authority and contractual relationships. There was a feeling that an instatiation of a P2 link between two NSCs was an agreement to work together in a certain way on UCPs.

The handling of a ticket between NSCs is bi-lateral. Should NSCs be certified to generate tickets? Should they be certified to accept tickets? Would one level of NSC be a "generate only" NSC while other NSCs could be "accept/generate" NSCs?

Every contact from a User (via the P1 protocol) must be logged and tracked by this system. The system must be conservative, it must not lose track of any calls (tickets) and it must reach closure on each ticket. What constitutes closure? All closures must be reported back to the User (via P1) and the User must be able to get status reports as the User requires (again via P1).

What are the minimum capabilities of an NSC? They should include:

- contact points (phone numbers, e-mail addresses, ...)
- hours of operation (when can the NSC be activated?)
- what do they do (ie, level of functionality)
- referrals (where do they refer UCPs via P2?)
- closure (they must be able to close open tickets via P1)

What is the role of UCP Central on routine UCPs? Should UCP Central get copies of all tickets from all NSCs? Should UCP Central be primarily mail based, as far as tracking tickets?

What is the nature of a ticket? The ticket must be structured such that it leads to a proper analysis of the problem. This implies a certain minimum of information. Can tickets be structured to include fields, as in a database? Guy Almes made the point that in talking about a distributed trouble ticket system, we are essentially trying to create a distributed database system. Perhaps we can glean some insight on how to structure P2 and create a coherent distributed trouble ticket system from distributed database design? Can we create a trouble ticket grammar? Should the trouble tickets be textual, so that they can be moved via mail, not requiring a database query language or other special protocol?

Educating End Users

Martyne Halgren of Cornell contributed a memo to the ucp list prior to this meeting, addressing issues regarding educating end-users, and described NETHELP and NETLEARN tools to accomplish the education process. Unfortunately, the entire session needed to be devoted to a discussion of the larger picture, and there was no time to delve into the end-user part of the model. Martyne's contribution was held for follow-up discussion at a later time.

Session Closure

The host outlined a minimum of three things that need work:

- NSC Requirements
- the P1 protocol
- the P2 protocol

The host twisted arms:

Matt Mathis agreed to work on NSC requirements, the P1, and the P2 protocols. Guy Almes agreed to work with Matt on the P2 issue. Dan Jordt also indicated willingness to contribute.

Follow-up discussion and postings of work in progress will be to the ucp list `ucp[-request]@nic.near.`:

ATTENDEES

Guy Almes	almes@rice.ed
Glee Cady	ghc@merit.edu
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Kent England	kwe@bu.edu
Metin Feridun	mferidun@bbn.com
Martyne Hallgren	martyne@tcgould.tn.cornell.edu
Gene Hastings	hastings@psc.edu
Tom Holodnik	tjh@andrew.cmu.edu
Wendy Huntoon	huntoon@a.psc.edu
Dan Jordt	danj@cac.washington.edu
Marilyn Martin	martin@cdnnet.ca
Matt Mathis	mathis@pele.psc.edu
Berlin Moore	prepnet@andrew.cmu.edu
Donald Morris	morris@ucar.edu
Dave O'leary	oleary@noc.sura.net
Lee Oattes	oattes@utcs.utoronto.ca
Mike Patton	map@lcs.mit.edu
Marc-Andre Pepin	pepin@crim.ca
Paul Pomes	paul-pomes@uiuc.edu
Karen Roubicek	roubicek@nnsc.nsf.net
Jim Sheridan	jsherida@ibm.com
Dana Sitzler	dds@merit.edu
Pat Smith	psmith@merit.edu
Mary Stahl	stahl@nisc.sri.com
Louis Steinberg	louiss@ibm.com
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Edward Vielmetti	emv@math.lsa.umich.edu
Carol Ward	cward@spot.colorado.edu
Aileen Yuan	aileen@gateway.mitre.org

3.2.7 User Documents (userdoc)

Charter

Chairperson:

Karen Roubicek, roubicek@nnsc.nsf.net
Tracy LaQuey,

Mailing Lists:

General Discussion: user-doc@nnsc.nsf.net
To Subscribe: user-doc-request@nnsc.nsf.net

Description of Working Group:

The USER-DOC Working Group will prepare a bibliography of on-line and hard copy documents/reference materials/training tools addressing general networking information and "how to use the Internet". (Target audience: those individuals who provide services to end users and end users themselves.)

- Identify and categorize useful documents/reference materials/training tools.
- Publish both an on-line and hard copy of this bibliography.
- Develop and implement procedures to maintain and update the bibliography. Identify the organization or individual(s) who will accept responsibility for this effort.
- As a part of the update process, identify new materials for inclusion into the active bibliography.
- Set up procedures for periodic review of the bibliograhy by USWG.

Goals and Milestones:

Done	Format for the bibliography will be decided as well as identification of "sources of information" (e.g., individuals, mailing lists, bulletins, etc.)
Done	Draft bibliography will be prepared
Mar 1990	Draft to be reviewed and installed in the Internet-Drafts Directory
May 1990	Bibliography submitted as a FYI RFC

CURRENT MEETING REPORT

Reported by Tracy LaQuey/ University of Texas

AGENDA

1. Review current version of bibliography
2. Discuss maintenance of bibliography, update procedures
3. Set deadline for corrections and submissions
4. Discuss Internet Draft procedure

MINUTES

The draft bibliography, which Aileen Yuan has been updating, was reviewed at the beginning of the meeting and some suggestions and corrections were made. Because of its length, the current version of the bibliography was made available electronically before the May IETF meeting. Copies of it were available for review, but were not distributed at this meeting. The following was discussed:

- There will be two versions generated from a single source (the refer database format) - a PostScript version and a plain text version. The refer format will also be made available so users can customize the bibliography for their needs.
- Because there has not been much response to the solicitations for material, we have not decided on a method for updating the bibliography. Five minute status reports will be given at future User Services Working Group meetings and decisions on whether or not we should update the bibliography will be made then. We will also report on any feedback or comments.
- It was decided that we should submit it as an Internet Draft as soon as possible. There will probably still be some missing pieces, and those will be filled in while it's in Internet Draft form. The Internet Draft should be available around the beginning of June.
- The draft will also be available on host nnsc.nsf.net.
- The RFC and FYI numbering scheme was discussed. The bibliography will be both an RFC and an FYI. It will be assigned a permanent FYI number. The RFC number will change if there are new versions.
- May 15 was set as the deadline for submissions and corrections to be sent to the USER-DOC mailing list. The USER-DOC list will be dissolved. Future messages regarding the bibliography should be sent to the USWG list.
- The USWG Distribution and Announcement Group (DAWG) will take care of advertising and distributing the bibliography.

Revised Publication Schedule

May 7 - 15	Working Group will tie up loose ends
May 15	Last day for submitting entries
June 1	Submit as an Internet Draft

ATTENDEES

Martyne Hallgren	martyne@tcgould.tn.cornell.edu
Tracy LaQuey	tracy@emx.utexas.edu
Berlin Moore	prepnet@andrew.cmu.edu
Donald Morris	morris@ucar.edu
Lee Oattes	oattes@utcs.utoronto.ca
Marc-Andre Pepin	pepin@crim.ca
Joyce Reynolds	jkrey@venera.isi.edu
Mike Roberts	roberts@educom.edu
Karen Roubicek	roubicek@nnsc.nsf.net
Tim Seaver	tas@mcnc.org
Dana Sitzler	dds@merit.edu
Mary Stahl	stahl@nisc.sri.com
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Aileen Yuan	aileen@gateway.mitre.org

3.2.8 User Services (uswg)

Charter

Chairperson:
Joyce Reynolds, jkrey@venera.isi.edu

Mailing Lists:
General Discussion: us-wg@nnsc.nsf.net
To Subscribe: us-wg-request@nnsc.nsf.net

Description of Working Group:

The User Services Working Group provides a regular forum for people interested in user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet. (Note that the actual projects themselves will be handled by separate groups, such as IETF WGs created to perform certain projects, or outside organizations such as SIGUCCS.

- Meet on a regular basis to consider projects designed to improve services to end-users. In general, projects should
  - clearly address user assistance needs;
  - produce an end-result (e.g. a document, a program plan, etc);
  - have a reasonably clear approach to achieving the end-result (with an estimated time for completion);
  - and not duplicate existing or previous efforts.
- Create WGs or other focus groups to carry out projects deemed worthy of pursuing.
- Provide a forum in which user services providers can discuss and identify common concerns.

Goals and Milestones:

Ongoing	This is an oversight group with continuing responsibilities.

CURRENT MEETING REPORT

Reported by Joyce Reynolds/ ISI and Glee Cady/ Merit

Announcements:

- New WG - SSPHWG - Site Security Policy Handbook
- SSPHWG's next meeting, Tuesday, June 12th, at USC/Information Sciences Institute
- Premier of the RFC FYI Series (RFC1150, FYI1)
- NOCtools Catalog Published (RFC1147, FYI2)
- Karen Bowers resignation from NRI as of April 30th
- The User Services Area report will be presented by Joyce Reynolds, need input from User-Doc, NISI Chairs.

Reports from User-Doc and NISI

Tracy LaQuey announced that the User-Doc Bibliography is ready for the Internet Draft Process. Final changes or amendments to the Bibliography have a deadline date of May 15th. After the Internet-Draft process then the RFC FYI publication, the User-Doc WG will terminate and go back into the USWG.

Dana Sitzler updated the USWG on NISI activities. NISI's first meeting focused on discussion of the old NISI charter and a survey of "where we are now" (i.e., a survey of existing informational types, retrieval mechanisms, and current NIC specialties and relationships). Specifically, how to get information to people. A draft will be sent to the NISI mailing list.

Distribution and Announcement Handbook

Bob Enger presented the DAWG document to the USWG.

- The immediate role of the IETF in broadly distributing information to the Internet community is to make use of communications avenues already developed by other organizations.
- The purpose of this handbook is to: 1) identify and provide specifics on various existing distribution resources, and to 2) consider possible long-term distribution methods.
- The intent is for this to be a handbook that can be used by all the IETF Working Groups to announce and/or distribute their documents as their charters dictate.
- There had been some question at the Tallahassee meeting whether DAWG should be in the USWG or NISI realm. It was decided that this handbook will stay with the USWG.

A suggested format of the DAWG template was drawn up:

To: (Audience)
From: (Provider)

How to obtain
Internet community
Other Audience

Discussion shifted to queries about, "Do we need to make this handbook an RFC??" The general concensus is that the DAWG handbook will be an IETF internal document only, specifically to help the IETF Area Directors and IETF WG Chairs "get the word out", beyond the normal distribution via RFC memos.

Quail Report

Gary Malkin was unable to attend, so Joyce Reynolds presented his report.

- Concern was expressed about authoritative answers.
- Questions shoud be generalized, so should answers.
- If answers are not definitve, the answer should not be given. We should bring the asker up to speed and then point him/her in the right direction for further information.
- There should be an update plan. Gary has planned to do so, at each IETF plenary.
- The Q/A draft document needs to be restructured.

The input from this discussion will be reported to Gary...further discussion will take place on the Quail mailing list.

"Intro Packages" - a new user electronic application

Continued discussion from last USWG meeting on what the information is going to be, what already exists, and what needs to be defined.

Additional research is needed. Martyne Hallgren, Karen Roubicek, and Joyce Reynolds will do further research and report at the next USWG session.

Next USWG meeting will be at UBC, where the USWG will continue discussion and research on:

- DAWG
- QUAIL

- Intro Packages

ATTENDEES

Alison Brown	alison@maverick@osc.edu
Ted Brunners	tob@thumper.bellcore.com
Jeff Ca'pente	jjc@unix.cis.pitt.edu
Wilson Dillaway	dillaway@sun.udel.edu
Greg Dobrich	dobrich@a.psc.edu
Robert Enger	enger@sccgate.scc.com
Martyne Hallgren	martyne@tcgould.tn.cornell.edu
Ole Jacobsen	ole@csli.stanford.edu
Tracy Laquey	tracy@emx.utexas.edu
Marilyn Martin	martin@cdnnet.ca
David Miller	dtm@mitre.org
Berlin Moore	prepnet@andrew.cmu.edu
Donald Morris	morris@ucar.edu
Marc-Andre Pepin	pepin@risq.net
Ron Roberts	roberts@jessica.stanford.edu
Karen Roubicek	roubicek@nnsc.nsf.net
Jim Sheridan	jsherida@ibm.com
Dana Sitzler	dds@merit.edu
Mary Stahl	stahl@nisc.sri.com
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
G W Cady	ghc@merit.edu ?
Carol Ward	cward@spot.colorado.edu
John Wobus	jmwobus@suvm.acs.syr.edu
Aileen Yuan	aileen@gateway.mitre.org

3.3 Internet Area

Director: Noel Chiappa

Area Summary
Reported by Greg Vaudreuil /CNRI

The Internet Area currently has 8 active working groups. Of these groups, the Connection Oriented IP, MTU Discovery, IP and Appletalk, IP over Switched Megabit Data Service, Point-to-Point Protocol Extensions, Router Discovery, and Router Requirements working groups met.

Both the Point-to-Point Protocol (PPP) documents have been submitted to the IAB for publication. The initial protocol document (RFC 1134) has been resubmitted to IAB for elevation to Draft Standard and the PPP Options has been submitted to the IAB for consideration as a Proposed Standard. The PPP Options document completes the specification and will be advanced with the initial document as a set. Language has been added to take into consideration the use of PPP over Frame Relay systems and to clarify security concerns.

The IP over Switched Megabit Data Service (SMDS) working group has posted an initial document to the Internet Drafts directory. The Connection IP working group presented the latest version of the ST protocol at this IETF plenary. A copy of the slides and a summary of the presentation are included later in this proceedings. They will release a specification shortly. The Performance and Congestion Control working group has risen from the dead to issue a new version of the Performance Internet-Draft.

3.3.1 IP over Appletalk (appleip)

<u>Charter</u>

Chairperson:

John Veizades, `veizades@apple.com`

Mailing Lists:

General Discussion: `apple-ip@apple.com`
To Subscribe: `apple-ip-request@apple.com`

Description of Working Group:

The Macintosh working group is chartered to facilitate the connection of Apple Macintoshes to IP internets and to address the issues of distributing AppleTalk services in an IP internet.

Goals and Milestones:

Feb 1991 Describe, in an RFC, the current set of protocols used to connect Macintoshes to IP internets.

Feb 1991 Define a MIB for the management of DDP/IP gateways.

CURRENT MEETING REPORT

Reported by John Veizades/ Apple

Minutes:

There was quite a bit of lively debate over the priorities of the working group, but all priorities involve the effective support of Macintosh computers on the Internet. AppleTalk over IP discussion:

The issues involving AppleTalk encapsulation over IP networks are these:

1. There's no standard for the existing state of IPTalk.
2. There are several areas where the current state of IPTalk might be improved.

The problems with IPTalk derive from the mismatch in pairing the IP/UDP layer with the AppleTalk DDP layer; a better match might be at the level of IP and DDP (i.e. encapsulate AppleTalk DDP packets just above the IP layer, beside UDP). However, this would come at the expense of making changes for every IP implementation in existence, which isn't feasible. There are also problems with the number of UDP ports a MacTCP machine can open, and the number of UDP ports the IPTalk server is required to maintain; an IPTalk machine (such as a UNIX machine running CAP) is required to listen on 256 UDP ports which are mapped to 256 AppleTalk DDP ports, while a MacTCP host can only maintain 64 UDP ports. Therefore, MacTCP machines can't fully interoperate with IPTalk machines.

AppleTalk might scale better over large networks if IP is used effectively as a transport.

Simplicity versus scale-ability. To what extent does support for large networks require extensive configuration from the maintainer? AppleTalk has always been constructed to be "plug-and-play," but that has introduced some problems with support over larger networks.

How well will AppleTalk Phase 2 be supported by IPTalk, if at all? IPTalk routing isn't documented anywhere except within the KIP code itself.

Documents describing Ed Moy's work (at UCB) were distributed. Since not everyone attending was familiar with the work, it was agreed to examine it, and follow up with it as a base for further work, as it seems to show considerable promise. Ed Moy's work not only attempts to document the existing state, but to propose a new IPTalk standard.

Ed Moy's report can be used as a starting point to address the issue where there is no documentation for the current state of IPTalk networking. It might also be used to address the problems with the current level of IPTalk networking. IP over AppleTalk Networks:

John Veizades (veizades@apple.com) presented an outline for a document to standardize the methods by which the IP is conducted over an AppleTalk network. The outline was generally accepted, and several areas were discussed.

An optional feature of the IP implementation on each Macintosh might be to send a packet to the IP address assignment agent to shutdown IP service. When a Macintosh completes a session and no longer requires an IP address, it may send a request to the gateway to free that address. If the feature is not implemented the address will age out of the assigning agents table of assigned addresses.

In discussion of the operation of higher layer protocols, two regimes were addressed: when the locally attached DDP-IP gateway is acting as an IP router, and when it's serving as an IP forwarding agent. If the DDP-IP gateway is serving as a router, it should comply with RFC-1009, the Router Requirements Specification. This would also require that the IP implementaion on all Macintoshes handle ICMP packets (of all varieties).

If the locally attached DDP-IP gateway is only forwarding IP packets, then "non-intuitive" things may occur when two IP-forwarders are connected to the same LocalTalk network, and connected to the same IP (sub)network. Proxy-ARP in this case leads to some confusion.

It was recommended that there should be no mention of DDP-ARP in the standards document.

The AppleTalk MIB:

The only reservations raised about the proposed MIB for AppleTalk were that the KIP section of the MIB had to refer to documented standard protocols (i.e. we need to document the KIP routing protocol), and that the buffer section had some FastPath-specific sections that might be better addressed in a vendor-specific MIB. In particular, the buffer section of the MIB might be geared more toward a FastPath than to any other product. Leaving information about buffer counts was agreed to be better left to a vendor-specific MIB.

Conclusions:

Several documents need to be drafted:

1. A specification for IP over AppleTalk (based on John Veizades' outline)
2. A specification for AppleTalk over IP (based on Ed Moy's report)
3. A further revision of the AppleTalk MIB (Steve Waldbusser's, with modifications)

ATTENDEES

Leo McLaughlin	ljm@twg.com
Rob Chandhok	chandok+@cs.cmu.edu
Bruce Crabill	bruce@umdd.umd.edu
Peter DiCamillo	cmsmaint@brownvm.brown.edu
Karen Frisa	karen@kinetics.com
Kanchei Loa	loa@sps.mot.com
Tom Holodnik	tjh@andrew.cmu.edu
Jonathan Wenocur	jhw@shiva.com
Mike Horowitz	mah@shiva.com
Frank Slaughter	fgs@shiva.com
Josh Littlefield	jash@cayman.com
Brad Parker	brad@cayman.com
Zbigniew Opalka	zopalka@bbn.com
Russ Hobby	rdhobby@ucdavis.edu
Van Jacobson	van@helios.ee.lbl.gov
Peter Vinsel	farcomp!pvc@apple.com
Terry Braun	tab@kinetics.com
Matthew Nocifore	matthew@durp.ocs.drexel.edu
Milo Medin	medin@nsipo.nasa.gov
David Kaufamn	dek@proteon.com
Steven Willis	swillis@wellfleet.com
Greg Satz	satz@cisco.com
Zaw-Sing Su	zsu@srz.com
John Veizades	veizades@apple.com

3.3.2 Connection IP (cip)

Charter

Chairperson:
Claudio Topolcic, topolcic@bbn.com

Mailing Lists:
General Discussion: cip@bbn.com
To Subscribe: cip-request@bbn.com

Description of Working Group:

This working group is looking at issues involved in connection-oriented (or stream- or flow-oriented) internet level protocols. The long term intent is to identify the issues involved, to understand them, to identify algorithms that address them, and to produce a specification for a protocol that incorporates what the working group has learned. To achieve this goal, the group is defining a two year collaborative research effort based on a common hardware and software base. This will include implementing different algorithms that address the issues involved and performing experiments to compare them. On a shorter time-line, ST is a stream-oriented protocol that is currently in use in the Internet. A short-term goal of this working group is to define a new specification for ST, called ST-2, inviting participation by any interested people. MCHIP and the Flow Protocol have also been discussed because they include relevant ideas.

Goals and Milestones:

Apr 1990	Produce a new specification of ST.
May 1990	Define common hardware and software platform.
Oct 1990	Implement hardware and software platform.
May 1991	Implement experimental modules and perform experiments.
May 1992	Produce a specification of a next generation connection oriented protocol.

CURRENT MEETING REPORT

Reported by Claudio Topolcic/ BBN

The CO-IP Working Group met at the May 1-4 IETF Meeting at Carnegie-Mellon University. During the Tuesday sessions we tried to pick up where we had left off in Florida State. We heard updates on DARTNet and the TWBNet. Tony Mazraani gave a progress report on the COIP kernel and a presentation on Washington University's work on Resource Management in Broadcast Lans. Work toward defining experiments for the DARTNet was hindered since not all the key people were present. We spent the balance of the sessions discussing the current draft of the ST-2 protocol specification.

Charlie Lynn had previously edited and distributed the current draft of the ST-2 protocol specification. He had also written up a number of issues that needed more thinking. The group discussed these issues and a few others that came up during the meetings.

A number of editorial comments to the draft were discussed. These included some minor restructuring to minimize repetition and increase clarity. More forward and backward pointers were suggested, as well as more examples. Numerous editing changes were suggested.

We discussed the relation between ST and IP. We decided to allow two forms of the ST header. The short form is as had previously been specified. A long form is structured like an IP header so that it can be processed by IP-only agents, and takes the place of the concept of IP encapsulation. The long form may also be used when IP security is required or to reduce either deliberate or accidental denial of service problems.

The issue of use of multicast lead to a lot of discussion. Ideally, we would like to be able for an ST agent to request that the local network dynamically create a multicast group for use by a stream, as its use could reduce the network bandwidth required to support the stream. Unfortunately, there does not seem to be much support for dynamic management of multicast addresses (how does a "user" dynamically request a multicast address at a given protocol layer, what agent(s) on a network robustly assign multicast addresses, how are the assigned addresses mapped into addresses for use above the network layer, e.g., IP multicast addresses, how are the assigned addresses reliably released/garbage collected, etc.).

It was felt that trying to create such a service was a challenging problem tangential to the work of the Working Group and should be delegated to some other group. The result was either to use replication instead of multicast, or to use static multicast groups. The problem with the former is wasted bandwidth, that with the latter is scaling – what were formerly separable problems (solvable by each stream independently) now become problems which must be solved in common by all streams on a network. HID negotiation is one example.

We discussed mechanisms by which changes could be made to established streams. For example, it may be desirable to allow a request to change a stream's bandwidth to allow a range of possible bandwidths. If the new target can only be added with decreased bandwidth, it would be desirable to decrease that stream's bandwidth, if that is allowed by the stream, when a new targe is added to the stream. These features causes some difficulties in coordinating the changes among the ST agents, as well as the applications, while maintaining the uninterrupted flow of data packets.

Other specific issues discussed included the following:

1. A Target cannot be an IP multicast group.
2. The ACCEPT message should be delayed until the HID negotiation has been completed.
3. We are not addressing the issues of spoofing (beyond the security features to be provided for IP by SDNS), intentional denial of service, or unintentional denial of service resulting from broken routes.
4. The structure of the "Group of Streams" specification.
5. Whether source routes would be strict, loose, strict in ST and loose in IP, or something else. This issue was not resolved.

ATTENDEES

Fred Bohle	fab@saturn.acc.com
Terry Braun	tob@kinetics.com
Stephen Casner	casner@isi.edu
Danny Cohen	cohen@isi.edu
Richard Fox	sytek!rfox@sun.com
Jonathan Goldick	goldick@b.psc.edv
Jack Hahn	hahn@umd5.umd.edu
Charles Lynn	clynn@bbn.com
Tony Mazraani	tonym@flora.wustl.edu
Zaw-Sing Su	zsu@sri.com
Ian Thomas	ian@chipcom.com
Claudio Topolcic	topolcic@bbn.com
Dave Wood	wood@gateway.mitre.org

3.3.3 IP MTU Discovery (mtudisc)

Charter

Chairperson:
Jeff Mogul, mogul@decwrl.dec.com

Mailing Lists:
General Discussion: mtudwg@decwrl.dec.com
To Subscribe: mtudwg-request@decwrl.dec.com

Description of Working Group:

The MTU Discovery Working Group is chartered to produce an RFC defining an official standard for an IP MTU Discovery Option. "MTU Discovery" is a process whereby an end-host discovers the smallest MTU along a path over which it is sending datagrams, with the aim of avoiding fragmentation.

Goals and Milestones:

Done — Decide if the proposal in RFC 1063 is sufficient, or if there are flaws to be corrected, or possible improvements to be made. Or, decide that it is unwise to create an official standard.

May 1990 — Unless the proposal in RFC 1063 is acceptable, write a new RFC describing a different approach.

Ongoing — Encourage the participation of gateway implementors, since the MTU discovery process affects the design and performance of IP gateways.

Done — Encourage sample implementations of end-host and gateway portions of MTU Discovery for popular software (BSD-derived kernels, primarily). (Encourage rapid implementation by major gateway vendors, since this option is relatively useless without widespread support.

CURRENT MEETING REPORT

Reported by Jeffrey Mogul/DEC

AGENDA

1. Report on current draft (Mogul/Deering)
2. Obtain consensus on approach
3. Focus on details
4. Schedule for standardization

MINUTES

This was the third meeting of the MTU Discovery Working Group.

Jeff Mogul started with a review of where we had been in the past 5 months, including all the failed approaches. He then presented the current proposal, originated by Steve Deering. Summary: send all packets with the DF bit set. Routers that cannot forward these packets return a slightly modified ICMP message that indicates the appropriate MTU. (Current routers return "0" in the field meant for this purpose.) The sending host revises its estimate of the Path MTU, and retransmits the dropped datagram.

There was no objection to the basic design. Some discussion ensued concerning details of the implementation and the relation between Path MTU discovery and transport protocol actions.

This design solves many problems with the previous designs, especially because it is compatible with existing hosts and routers, and is quite simple to implement.

ACTION ITEMS

1. Jeff Mogul and Steve Deering will change some details in the current draft and submit it as an Internet Draft.
2. The Router Requirements Working Group should be notified that we no longer care how fragmentation is done, since we do not rely on the size of fragments. We will also recommend that support for MTU Discovery be a requirement.

SCHEDULE

We expect never to meet again. Progress towards standardization will go as fast as possible, unless serious objections are raised.

ATTENDEES

Fred Bohle	fab@saturn.acc.com
Steve Bruniges	
David Burdelski	daveb@ftp.com
Duane Butler	dmb@network.com
Andrew Cherenson	arc@sgi.com
Noel Chiappa	jnc@ptt.lcs.mit.edu
Steve Deering	deering@pescadero.stanford.edu
Dave Forster	forster@marvin.enet.dec.com
Rich Fox	sytek!rfox@sun.com
Karen Frisa	karen@kinetics.com
Steve Hubert	hubert@cac.washington.edu
Van Jacobson	van@helios.ee.lbl.gov
Stev Knowles	stev@ftp.com
Yoni Malachi	yoni@cs.stanford.edu
Keith McCloghrie	sytek!kzm@hplabs.hp.com
Leo J. McLauglin III	ljm@twg.com
Jeff Mogul	mogul@decwrl.dec.com
Drew Perkins	ddp@andrew.cmu.edu
John Moy	jmoy@proteon.com
Stephanie Price	cmcvax!price@hub.ucsb.edu
Mike Reilly	reilly@nsl.dec.com
Tim Seaver	tas@mcnc.org
Frank Slaughter	fgs@shiva.com
Richard Smith	smiddy@pluto.dss.com
Brad Strand	bstrand@cray.com
Cal Thixton	cthixton@next.com
John Veizades	veizades@apple.com
Jonathan Wenocur	jhw@shiva.com

PATH MTU DISCOVERY

AGENDA

CURRENT STATUS

IS THIS ONE THE RIGHT APPROACH?

MISSING PIECES

SCHEDULE FOR STANDARDIZATION

CURRENT STATUS

YET ANOTHER DESIGN

DRAFT RFC (FIRST DRAFT)

HOST IMPLEMENTATION FOR 4.3BSD

? IMPLEMENTATION FOR 4.4BSD

ROUTER IMPLEMENTATIONS?

REJECTED SCHEMES (SO FAR...)

"576 RULE"
TOO CRUDE

RFC 1063
REQUIRES IP OPTION, ROUTER SUPPORT
HARD TO SCHEDULE REPLY OPTION

MCCLOGHRIE/FOX DRAFT
REQUIRES IP OPTION (ROUTER SUPPORT WOULD BE NICE)
COMPLEX HOST IMPLEMENTATION
LOTS OF PARAMETERS

"RF BIT"
MANY ROUTERS DROP THESE PKTS.
POLITICS OF HEADER BIT
REQUIRES ROUTER UPDATES BEFORE IT BECOMES USEFUL

OTHER ISSUES:
PC-IP BUGS
COST/BENEFIT

"DF-BIT" SCHEME

INITIAL PMTU ESTIMATE = FIRST-HOP MTU

TRANSPORT LAYER USING PMTU DISCOVERY SENDS ALL PACKETS WITH DF BIT SET

LOOPS UNTIL RIGHT

ROUTERS SEND ICMP "DESTINATION UNREACHABLE / FRAGMENTATION NEEDED AND DF SET" MESSAGE ("DATAGRAM TOO BIG") WITH USABLE MTU IN CURRENTLY UNUSED FIELD

ICMP MESSAGE CAUSES PMTU ESTIMATE DECREASE TO VALUE RETURNED BY ROUTER

TRANSPORT LAYER USES NEW PMTU VALUE TO SET PKT SIZE (E.G., MSS)

ESTIMATE DISCARDED EVERY ~10 MINUTES IN CASE ROUTE-CHANGE ALLOWS BIGGER PACKETS

COMPATIBILITY WITH EXISTING ROUTERS

CURRENT SPECS REQUIRE
→ "UNUSED" FIELD = 0
→ RETURNS ORIGINAL IP LENGTH *

ICMP MESSAGE CAUSES REDUCTION/CHANGE IN PMTU ESTIMATE

POSSIBLE ESTIMATION ALGORITHMS:

- GEOMETRIC SERIES ($E = E(1-\alpha)$)
  - INACCURATE RESULTS
  - SLOW IN WORST CASE
- BINARY SEARCH
  - NOT FAST IN BEST/COMMON CASES
  - HARD TO INCREASE ESTIMATES
- TABLE OF LIKELY VALUES
  - EXISTING MTUS CLUSTER NICELY
  - ≈ BINARY SEARCH (# OF STEPS) IN WORST CASE
  - O(1) STEP IN COMMON CASES
  - TRIVIAL IMPLEMENTATION
  - REQUIRES COORDINATION OF NEW MTU CHOICES

TRANSPORT ACTIONS

"DATAGRAM TOO BIG" MESSAGE MEANS PKT WAS DROPPED

RETRANSMISSION STRATEGY AFFECTS PERFORMANCE

SHOULD NOT INFLUENCE RTT ESTIMATE

POSSIBLE PROTOCOL MOD:
IF RF+DF BOTH SET
→ SEND ICMP
→ BUT ALSO FORWARD DATAGRAM

ROUTER DISCOVERY

- STATUS (RFC, IMPLEMENTATIONS)
- PREF. LEVEL FIELD
  - NEEDED?
  - HOW ENCODED?
  - HOW SET?
- MULTICAST VS. UNICAST REPLIES TO QUERIES
- DALLYING
  - WHEN?
  - HOW MUCH
- REPORTING RATE
- BLACK HOLE DETECTION

3.3.4 IP over FDDI (fddi)

Charter

Chairperson:
Dave Katz, dkatz@merit.edu

Mailing Lists:
General Discussion: FDDI@merit.edu
To Subscribe: FDDI-request@merit.edu

Description of Working Group:

The IP over FDDI Working Group is chartered to create Internet Standards for the use of the Internet Protocol and related protocols on the Fiber Distributed Data Interface (FDDI) medium. This protocol will provide support for the wide variety of FDDI configurations (e.g., dual MAC stations) in such a way as to not constrain their application, while maintaining the architectural philosophy of the Internet protocol suite. The group will maintain liason with other interested parties (e.g., ANSI ASC X3T9.5) to ensure technical alignment with other standards. This group is specifically not chartered to provide solutions to mixed media bridging problems.

Goals and Milestones:

May 1990 Write a document specifying the use of IP on a single MAC FDDI station.

Aug 1990 Write a document specifying the use of IP on dual MAC FDDI stations.

3.3.5 IP over Switched Megabit Data Service (smds)

Charter

Chairperson:

George Clapp, meritec!clapp@bellcore.bellcore.com
Mike Fidler, ts0026@ohstvma.ircc.ohio-state.edu

Mailing Lists:

General Discussion: smds@nri.reston.va.us
To Subscribe: smds-request@nri.reston.va.us

Description of Working Group:

The SMDS Working Group is chartered to investigate and to specify the manner in which the Internet and the newly defined public network service, Switched Multi-megabit Data Service, will interact. The group will discuss topics such as addressing, address resolution, network management, and routing.

Goals and Milestones:

TBD	Specify clearly an efficient interworking between the Internet and SMDS.

CURRENT MEETING REPORT

Reported by George Clapp/ Ameritech

The IP over SMDS Working Group met for three half-day sessions. During the first session, George Clapp presented a detailed tutorial on SMDS, the IEEE 802.6 MAN, and Broadcast ISDN. Copies of the slides are included in the Proceedings.

The second session was devoted to discussion of how ARP will be supported by SMDS, in which the IP address is the protocol address and the SMDS address is the hardware address. Discussion made it clear that there were a number of possible extensions which might be made to the network model of SMDS. These extensions increased the flexibility of the model but appeared to complicate the support of ARP. In the interest of simplifying the problem and generating an RFC quickly, a constrained set of conditions were listed.

1. Everyone in a closed group is in the same IP network/subnet.
2. Everyone else is accessed via a router.
3. The IP network/subnet will be bound to a single SMDS Group Address.
4. The broadcast MAC address is the SMDS group address. (This must be configured for each individual station in the closed group.)

An additional assumption for the baseline set of conditions was that IP would not be broken.

George Clapp volunteered to write a first draft of an RFC using the model which may be labeled the Virtual Private Network (VPN) model.

An alternative model is a "global" one in which it is assumed that any SMDS device may talk directly with any other SMDS device. Consistency with IP requires that all SMDS devices must belong to the same IP network/subnet.

ARP would be supported in the small VPN by multicasting the ARP packets to each member of the VPN. ARP would be supported in the global model by multicasting the ARP packets to a set of servers.

An additional parameter of the SMDS model was the type of devices which would be attached to the network, either all hosts, all routers, or a mixture of the two. It was pointed out that a network consisting of all hosts would be an isolated IP subnet.

A number of comments were made by the group during discussion:

- Would it be possible to use an algorithm to derive the SDMS group address from the IP network address?

- How would reverse ARP work?
- The MTU for SMDS is 9188 octets.
- How would bridging work across an SMDS network?
- The WG should look to the RFC describing IP over FDDI for guidance.
- A device with a single SMDS interface may belong to multiple IP networks/subnets.

Network Management Discussion.

A number of questions were raised during discussion of network management. It was suggested that the protocol would probably be SNMP, but that the order of business for the group should be to list, categorize, and prioritize issues to build the management model and then pick a protocol. The following issues were listed:

- Performance of the physical layer. It was pointed out that there was a separate Working Group on this topic and that the group should refer the task of building a physical MIB to them. The T1 MIB could be used as a reference. Some of the group (Bellcore folks?) indicated they would coordinate this.
- Maintenance of statistics, such as byte counts, packet counts, and CRC error counts. The question arose whether a device which kept these statistics would query the SMDS network for similar statistics for comparison.
- The cost structure would have an important impact on determining which statistics should be maintained.
- Provisioning and the "subscriber service profile". The management of the SMDS group addresses arose in the category. Presently, the SMDS group addresses are statically defined but the ability to dynamically add or delete group members is desirable. Another aspect of the service profile is access class. Statistics should be maintained of the number of packets dropped due to exceeding the access class bandwidth.

At the end of the last session, the following items were noted for further action:

- Dave Piscitello offered to publish a list of candidate objects for network management. He asked the WG to respond by email as to the importance of each object.
- George Clapp will write a first draft of an RFC defining the operation of ARP over an SMDS VPN.
- The group will initiate contact with the Transmission MIB WG concerning the definition of the physical layer MIB. (Who is doing this? - mlf)

The group then adjourned until the next IETF meeting at the University of British Columbia, Canada.

ATTENDEES

Z Bigniew Opalka	zopalka@bbn.com
Dave Borman	dab@cray.com
Lan Bosack	bosack@mathom.cisco.com
Theodore Brunner	tob@thumper.bellcore.com
Duane Butler	dmb@network.com
Allen Cargille	cargille@wisc.edu
Jeffrey Case	case@utkux1.utk.edu
Samir Chatterjee	samir@nynexst.com
Caroline Cranfill	rcc@blsouth.com
Dave Crocker	dcrocker@nsl.dec.com
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Robert Enger	enger@seka.scc.com?
Roger Fajman	raf@cu.nih.gov
Mike Fidler	mlf+@osu.edu
Richard Fox	sytek!rfox@sun.com
Patrick Glasso	pjg@sei.cmu.edu
Phill Gross	pgross@nri.reston.va.us
Robert Hagens	hagens@cs.wisc.edu
Jack Hahn	hahn@umd5.umd.edu
Tony Hain	hain@nmfecc.arpa
Bob Hinden	hinden@bbn.com
Russell Hobby	rdhobby@ucdavis.edu
Keith Hogan	keith%penril@uunet.uu.net
Kathy Huber	khuber@bbn.com
Patrick J. Glasso	pjg@sei.cmu.edu
Van Jacobson	van@helios.ee.lbl.gov
Tony Lauck	lauck@dsmail.dec.com
Joe Lawrence	jcl@sabre.bellcore.com
James Leighton	jfl@nmfecc.arpa
Mark Leon	leon@nsipo.arc.nasa.gov
George Marshall	george@adapt.net.com
Tony Mazraani	tonym@flora.wustl.edu
Donald Merritt	don@brl.mil
Cyndi Mills	cmills@bbn.com
Dave O'leary	oleary@rec.sura.net
Park Parker	paul.parker@cs.cmv.edv
David Piscitello	dave@sabre.bellcore.com
Joel Replogle	replogle@ncsa.uiuc.edu
Mike Roberts	roberts@educom.edu
Ron Roberts	roberts@jessica.stanford.edu

Robert Sansom	sansom@cs.cmu.edu
Jim Showalter	gamma@mintaka.dca.mil
Frank Slaughter	fgs@shiva.com
Bob Stafford	stafford@fac.cis.temple.edu
Zaw-Sing Su	zsu@tsca.istc.sri.com
Thomas Swazuk	swazuk@fac.cis.temple.edu
Gregory Vaudreuil	gvaudre@nri.reston.va.us
Steve Willis	swillis@wellfleet.com
Dan Wintringham	danw@igloo.osc.edu
Dave Wood	wood@gabeway.mitre.org
Chin Yuan	cyuan@srv.pacbell.com

Switched Multi-megabit Data Service, Metropolitan Area Networks, & Broadband ISDN

George H. Clapp
Ameritech Services, Inc.
Science and Technology
Gould Center, Building 40
2850 Golf Road
Rolling Meadows, IL 60008-4014
708-806-8318
fax: 708-806-8292
email: meritec!clapp@bellcore.bellcore.com
clapp@maui.cs.ucla.edu

Topics

- Standards activities
- Network architectures
- MAN Services and Functional Blocks
- Distributed Queue Dual Bus (DQDB) Protocol
  - Operation
  - Structures
  - Performance
  - Physical Layer Convergence Protocol

Broadband Services

Conversational Services
- Video Telephony
- Broadband Teleconference
- Video Surveillance
- High-Speed Telefax
- LAN Interconnect
- Host-to-Host
- Real Time Control

Distribution Services
- Existing Quality TV
- Extended Quality TV
- High Definition TV
- Pay TV
- Multi-Lingual TV
- Audio Distribution
- Full Channel Videotex

Messaging
- Video Mail
- Document Mail

Retrieval Services
- Broadband Videotex
- Remote Education
- Entertainment
- Document Retrieval
- High Res Image Retrieval
- "Mixed" Documents
- Electronic Newspapers
- Telesoftware

What is Switched Multi-megabit Data Service (SMDS)?

- Technical Advisory (TA-TSY-000772) released by Bell Communications Research (Bellcore).
- High-speed, connectionless, public, packet switching service which will extend LAN-like performance beyond the subscriber's premises.
- Transmission rates are DS3 (44.736 Mb/s line signaling rate with 44.209 Mb/s payload) and DS1 (1.544 Mb/s line signaling rate with 1.536 Mb/s payload).
- Issue 1 was released in February of 1988; Issue 2 in March of 1989.

What is an IEEE 802.6 MAN?

- Standardization of the ***Distributed Queue Dual Bus (DQDB)*** Medium Access Control (MAC) algorithm, which resolves contention to a shared medium, broadcast network.
- Extension of a Local Area Network in speed, distance, and number of users.
- Integrated transport of high speed data, voice, and compressed video.
- ≥50 km. in diameter.
- Primary service is high speed connectionless data transport and switching.
- Initial transmission line signaling rate will DS3 (45 Mb/s) with extension to SONET (Synchronous Optical NETwork) rates (155.52 Mb/s).
- Public Network.
- Standardization work began in April of 1981.

What is Broadband ISDN?

- Extension of ISDN in speeds and services
- Integrated transport of high speed data, voice, and video.
- Motivated by...
  - Fiber optic technology.
  - Vehicle for the distribution of entertainment video.
  - Vehicle for high speed data transport and switching.
- Tentative line signaling rates of 155.520 Mb/s and 622.080 Mb/s.
- Standardization work began in January of 1985.

Bellcore Technical Advisory Review Process

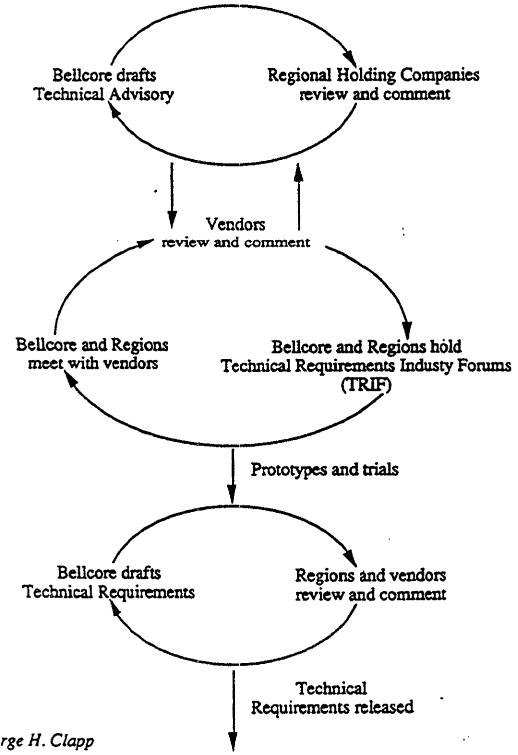
Bellcore drafts Technical Advisory
Regional Holding Companies review and comment
Vendors review and comment
Bellcore and Regions meet with vendors
Bellcore and Regions hold Technical Requirements Industy Forums (TRIF)
Prototypes and trials
Bellcore drafts Technical Requirements
Regions and vendors review and comment
Technical Requirements released

Standards Activities
Metropolitan Area Networks

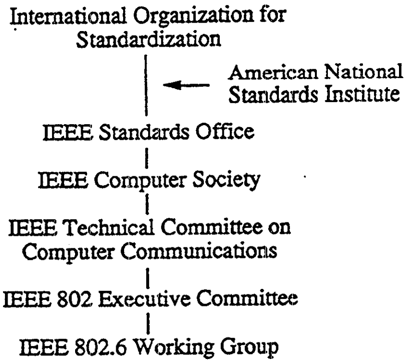
International Organization for Standardization
American National Standards Institute
IEEE Standards Office
IEEE Computer Society
IEEE Technical Committee on Computer Communications
IEEE 802 Executive Committee
IEEE 802.6 Working Group

MAN Architecture
Stand-alone 802.6 Subnetwork

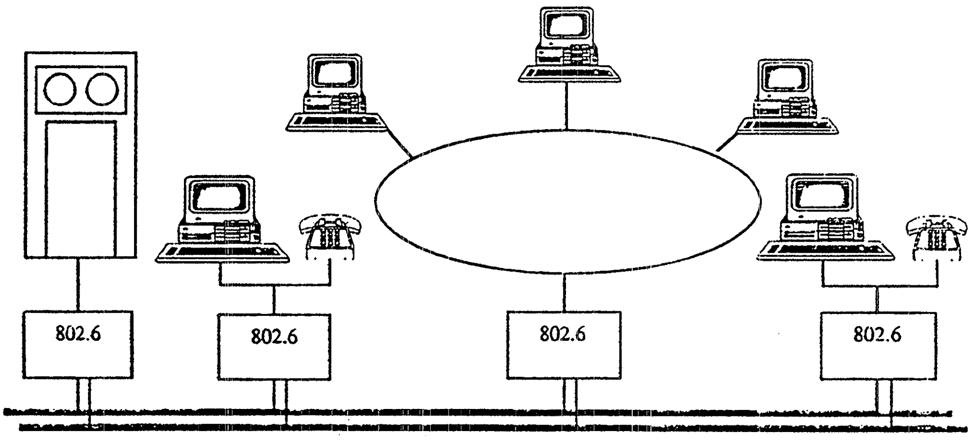

Insufficient Capacity

Standards Activities
Broadband ISDN

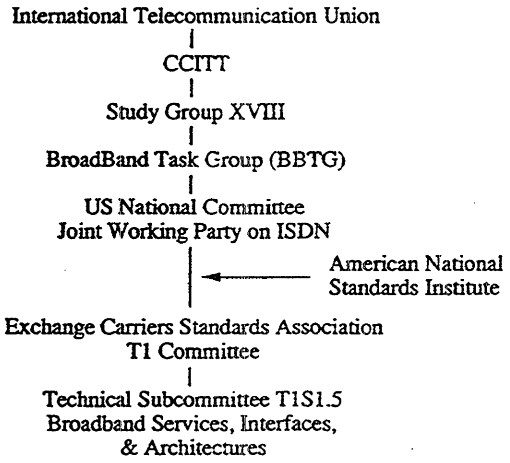

International Telecommunication Union

CCITT

Study Group XVIII

BroadBand Task Group (BBTG)

US National Committee
Joint Working Party on ISDN

American National
Standards Institute

Exchange Carriers Standards Association
T1 Committee

Technical Subcommittee T1S1.5
Broadband Services, Interfaces,
& Architectures

MAN Architecture
MANs Built as a Hierarchy of LANs

Limitations
- Capacity
- Performance
- Maintenance

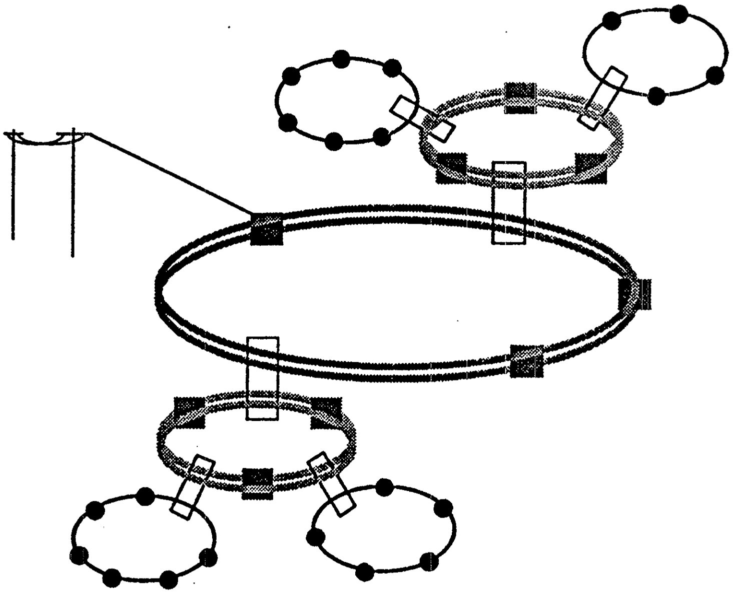

MAN Architecture
Introduction of a Central Switch

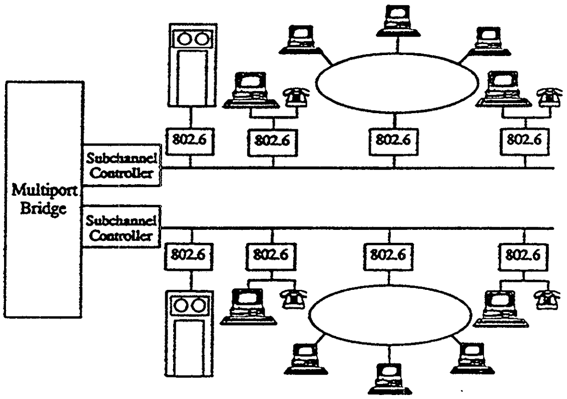

Target Broadband Data Architecture

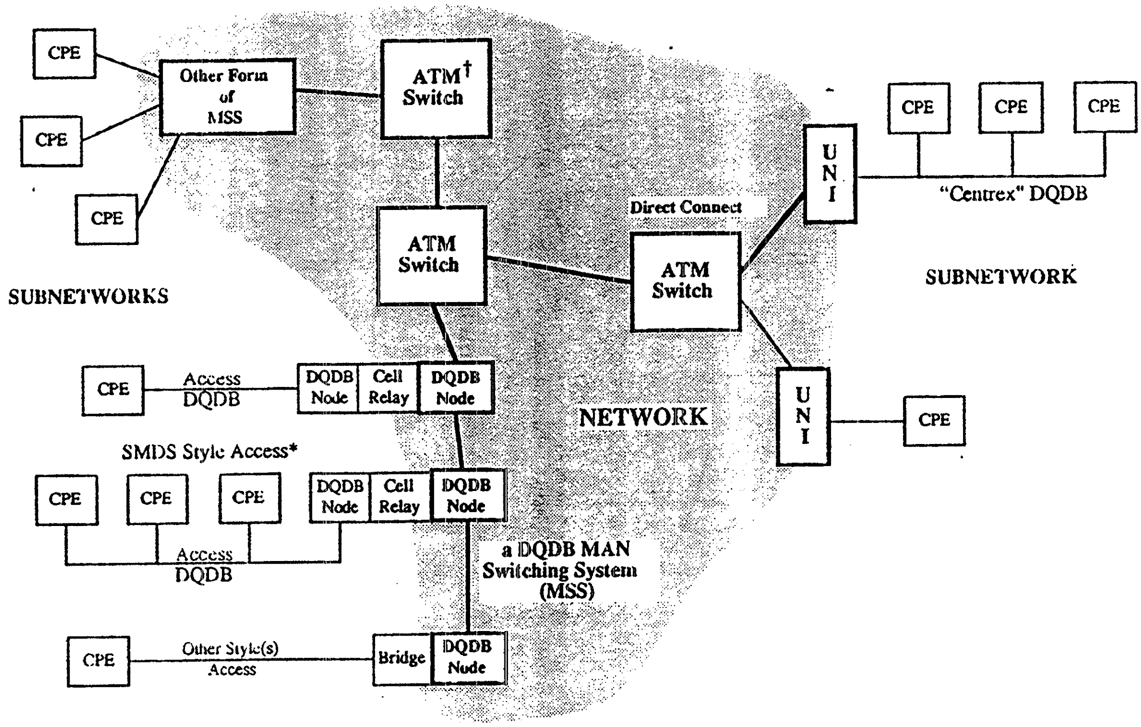

* Per Current Bellcore TA-772

† ATM Switches can initially be DQDB Bridges

George H. Clapp

AMERITECH SERVICES

Switched Multi-megabit Data Service (SMDS)

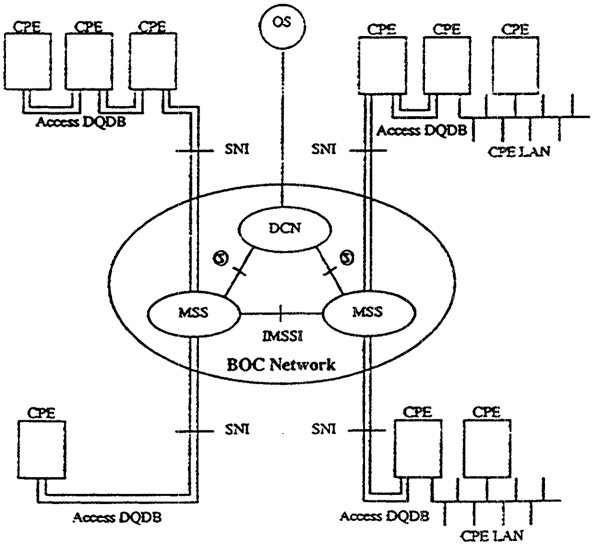

SNI:	Subscriber Network Interface	DCN:	Data Communications Network
MSS:	MAN Switching System	CPE:	Customer Premises Equipment
IMSSI:	Inter-MAN Switching System Interface	CPE LAN:	CPE Local Area Network
Ⓢ:	Generic Interface for Operations	DQDB:	Distributed Queue Dual Bus
OS:	Operations System		

George H. Clapp

AMERITECH SERVICES

Broadband ISDN Interface:
Asynchronous Transfer Mode (ATM)

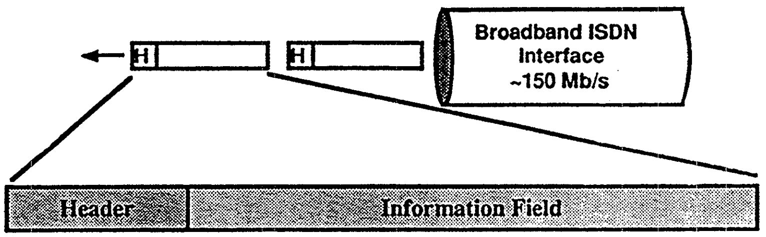

Characteristics:

- Common packet-like capability. capable of supporting all services.
- Consists of a streams of "cells" with fixed-length headers and information fields.
- Individual conversations identified by a virtual channel identifier in the headers and not by the location of the cell in a frame.

☞ The target architecture of Broadband ISDN will be based on ATM.

Why ATM?

Flexibility for the End User:

- Ability to realize arbitrary size circuits.
- Allows any combination of synchronous and asynchronous traffic including multimedia services.
- Provides dynamic allocation of bandwidth on demand.

Flexibility for the Network Operator:

- Ability to mix different traffic types in the same network.
- Facilitates switching/transmission integration.
- Adapts to changing customer bandwidth requirements.
- Could allow operation without synchronous clock hierarchies.
- Simplified network architectures.
- Easy add/drop of bandwidth.
- Efficient use of bandwidth.

IEEE Project 802

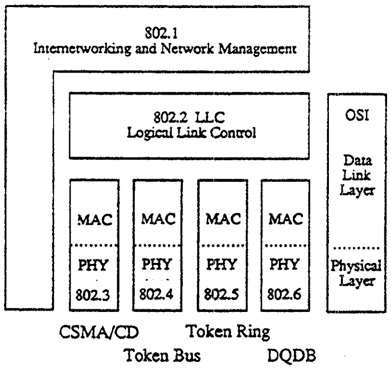

Also...

802.7: Broadband Technical Advisory Group

802.8: Fiber Optic Technical Advisory Group

802.9: Integrated Voice & Data (IVD) Working Group

802.10: Standard for Interoperable LAN Security (SILS) Working Group

George H. Clapp

Services of the IEEE 802.6 MAN

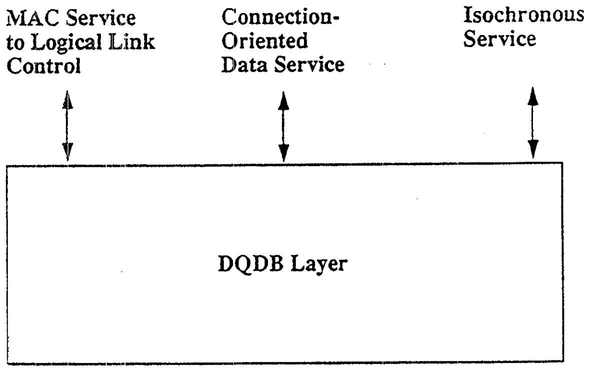
MAC Service to Logical Link Control
Connection-Oriented Data Service
Isochronous Service
DQDB Layer

George H. Clapp

DQDB
Functional Block Diagram

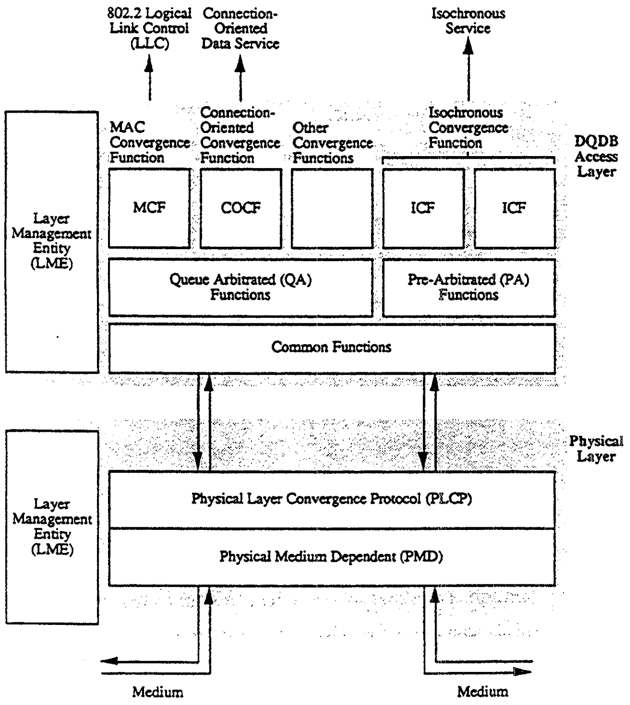
802.2 Logical Link Control (LLC)
Connection-Oriented Data Service
Isochronous Service
Layer Management Entity (LME)
MAC Convergence Function
Connection-Oriented Convergence Function
Other Convergence Functions
Isochronous Convergence Function
DQDB Access Layer
MCF
COCF
ICF
ICF
Queue Arbitrated (QA) Functions
Pre-Arbitrated (PA) Functions
Common Functions
Physical Layer
Layer Management Entity (LME)
Physical Layer Convergence Protocol (PLCP)
Physical Medium Dependent (PMD)
Medium
Medium

George H. Clapp

IEEE 802.6 Dual Bus
"Open" Bus

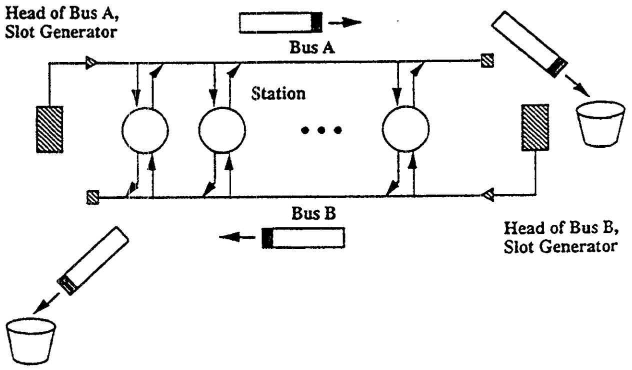

- Two uni-directional, counter-flowing buses. Capacity of the dual bus is twice the capacity of a single bus.
- Stations have read/write access to both buses.
- Slots transmitted by slot generator "fall off" the end of the bus.

George H. Clapp

IEEE 802.6 Dual Bus
"Looped" Bus

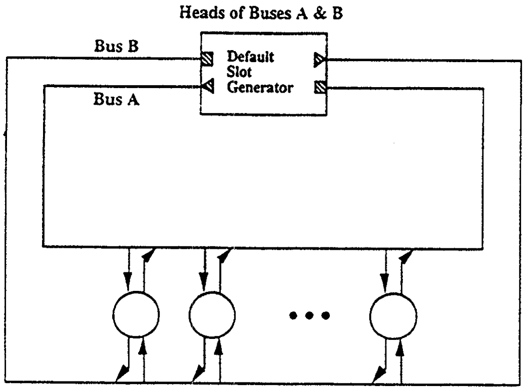

- Capable of reconfiguration in the event of a single link failure.
- Share "head of bus" functionality.

DQDB Subnetwork
Reconfiguration

Physical breaks in transmission links are healed by repositioning the natural break in the loop.

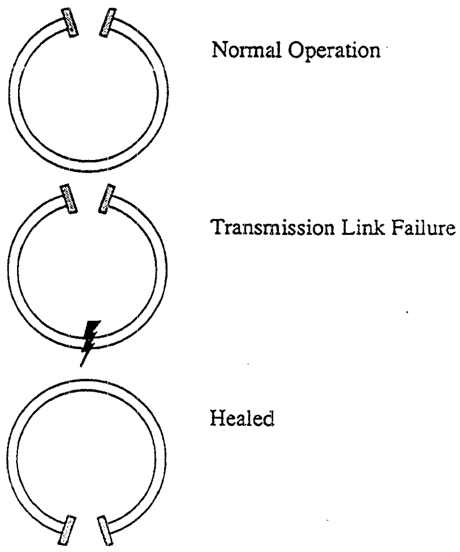

Station Architecture

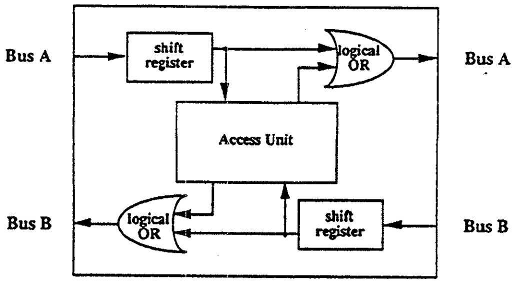

- Reading is done prior to writing.
- Writing is a logical OR function (at the DQDB access layer, not at the physical layer).

Distributed Queueing Access Protocol
Operation for Access to Bus A

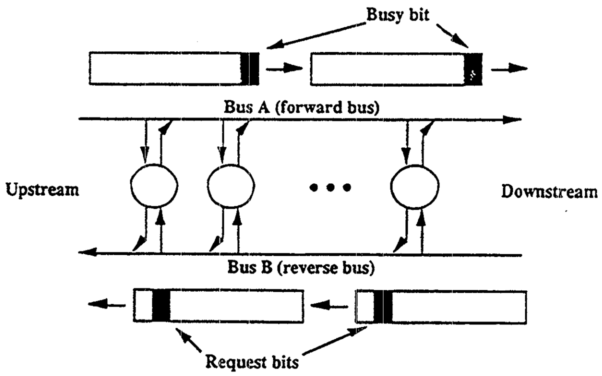

- The Slot Generator at the Head of Bus transmits fixed length segments.
- Busy bit indicates whether the segment is occupied.
- Request bits on the reverse bus indicate whether "downstream" stations wish to transmit on the forward bus.

Distributed Queue Access Protocol
Idle Station

Maintain a request counter (REQ_CNTR) to keep track of the segments queued downstream.

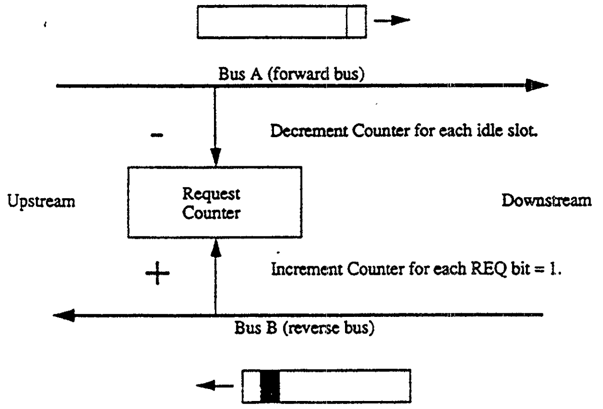
Bus A (forward bus)
-
Decrement Counter for each idle slot.
Upstream
Request Counter
Downstream
+
Increment Counter for each REQ bit = 1.
Bus B (reverse bus)

Distributed Queue Access Protocol
Queueing a Segment on Bus A

Operations:

1. Initiate transmission of REQ on reverse bus.
2. Transfer contents of REQ_CNTR to a "countdown" counter (CD_CNTR).
3. Set REQ_CNTR to zero.

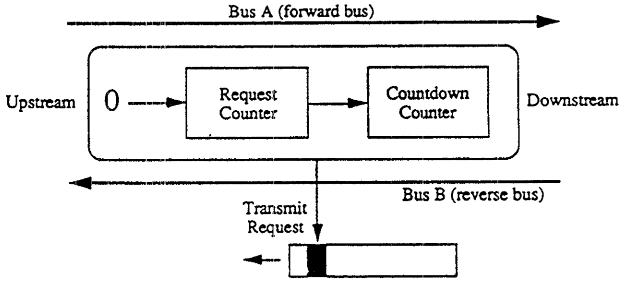

Distributed Queue Access Protocol
Transmitting a Segment on Bus A

Operations:

1. Decrement CD_CNTR for each idle slot.
2. Increment REQ_CNTR for each new request.
3. When CD_CNTR = 0, transmit segment in first idle slot.
4. Return to idle state.

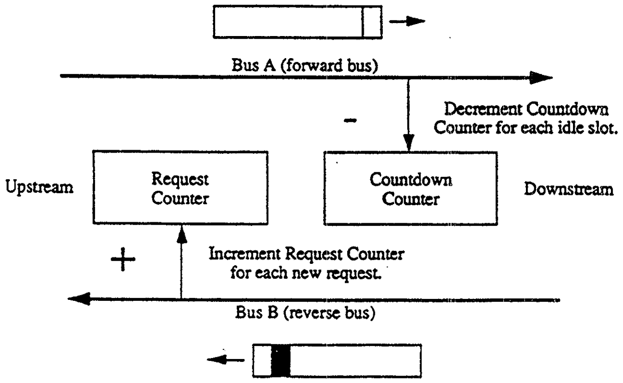
Bus A (forward bus)
-
Decrement Countdown Counter for each idle slot.
Upstream
Request Counter
Countdown Counter
Downstream
+
Increment Request Counter for each new request.
Bus B (reverse bus)

Distributed Queue Access Protocol
Example

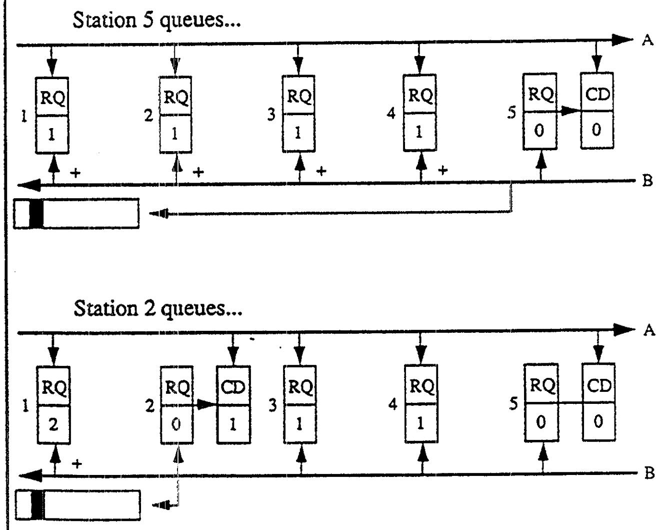

Distributed Queue Access Protocol
Example

Station 3 queues...

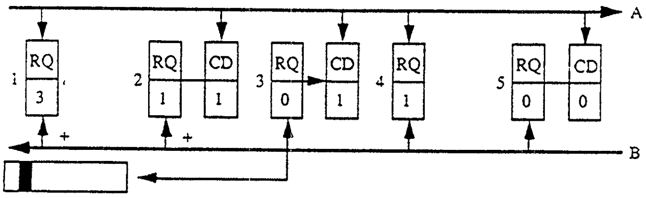

Stations 5, then 2, then 3 are queued...

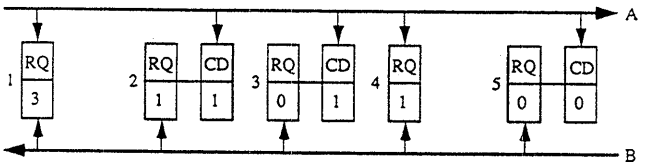

Distributed Queue Access Protocol
Example

Station 5 gains access...

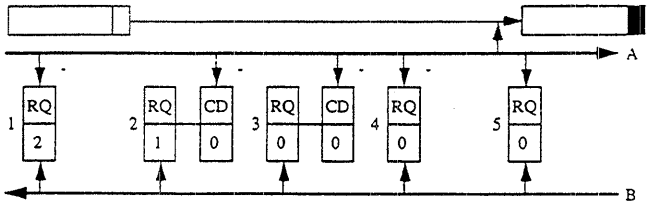

Station 2 gains access...

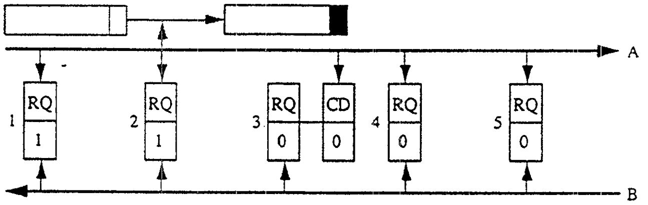

Distributed Queue Access Protocol
Example

Station 3 gains access...

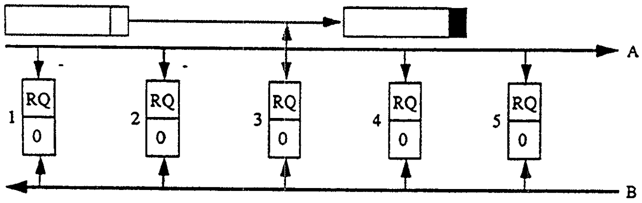

DQDB
Queue Arbitrated State Transition Diagram
(for Bus x at Priority I)

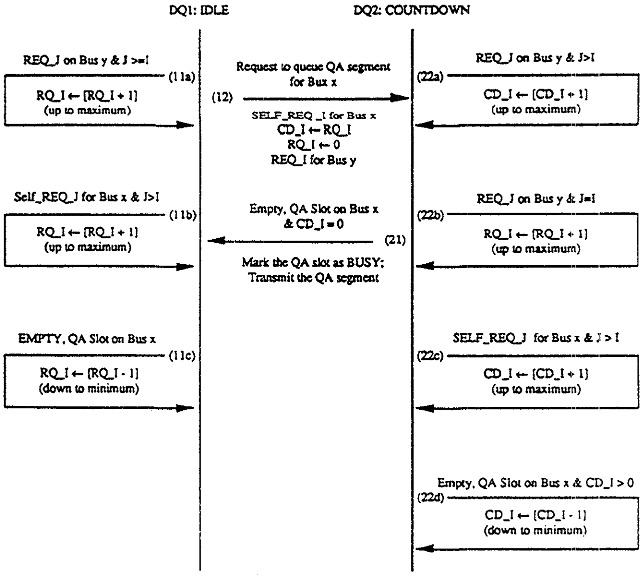

DQ1: IDLE

DQ2: COUNTDOWN

(11a) REQ_J on Bus y & J >=I
RQ_I ← [RQ_I + 1]
(up to maximum)

(11b) Self_REQ_J for Bus x & J>I
RQ_I ← [RQ_I + 1]
(up to maximum)

(11c) EMPTY, QA Slot on Bus x
RQ_I ← [RQ_I - 1]
(down to minimum)

(12) Request to queue QA segment for Bux x
SELF_REQ _I for Bus x
CD_I ← RQ_I
RQ_I ← 0
REQ_I for Bus y

(21) Empty, QA Slot on Bus x & CD_I = 0
Mark the QA slot as BUSY;
Transmit the QA segment

(22a) REQ_J on Bus y & J>I
CD_I ← [CD_I + 1]
(up to maximum)

(22b) REQ_J on Bus y & J=I
RQ_I ← [RQ_I + 1]
(up to maximum)

(22c) SELF_REQ_J for Bus x & J > I
CD_I ← [CD_I + 1]
(up to maximum)

(22d) Empty, QA Slot on Bus x & CD_I > 0
CD_I ← [CD_I - 1]
(down to minimum)

Fairness[†]

Under certain conditions (long network, high bandwidth, heavy traffic), there is the potential for unfairness.

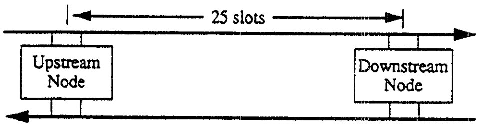

If both stations are fully loaded and...
- if the upstream node starts first,
  - upstream node gets 98%,
  - downstream node gets 2%;
- if the downstream node starts first,
  - upstream node gets 12%,
  - downstream node gets 88%.

[†] Viewgraph content by Ellen L. Hahne, AT&T Bell Laboratories, Murry Hill, NJ.
Improving DQDB Fairness, E. L. Hahne, N. F. Maxemchuk, A. K. Choudhury, IEEE Document 802.6-89/52.
Improving the Fairness of Distributed-Queue-Dual-Bus Networks, E. L. Hahne, A. K. Choudhury, N. F. Maxemchuk, Submitted to Infocom '90.

Fairness
Solution

Do not saturate the bandwidth. Each station takes ≤ 8/9 of the spare capacity. For every eight segments transmitted, station declines one opportunity to transmit a segment.

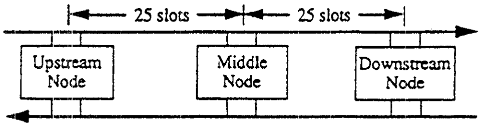

Throughput over 100 slot intervals

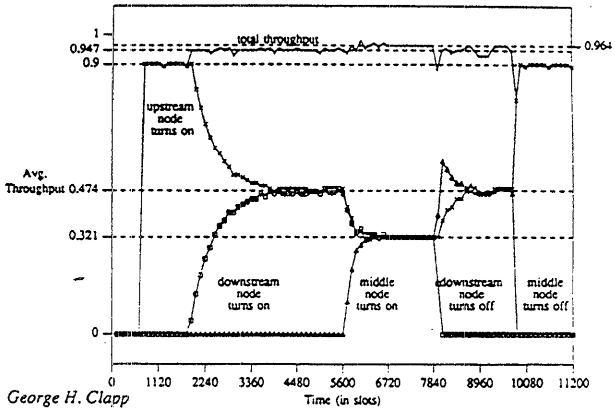

Transmission of a Packet

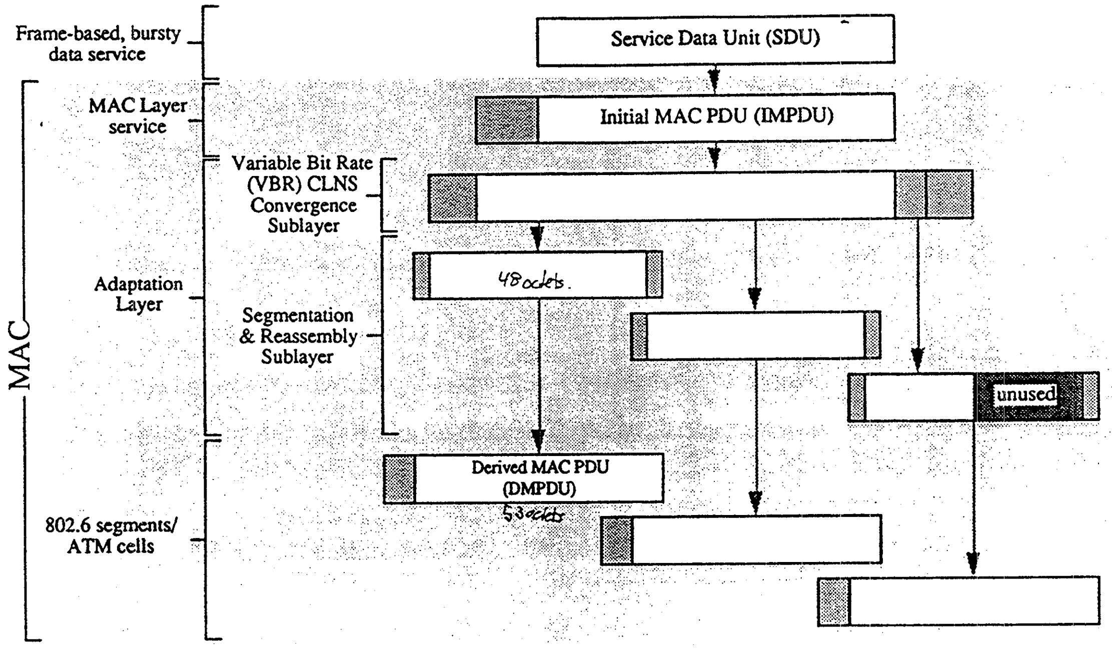
Frame-based, bursty data service
MAC Layer service
Variable Bit Rate (VBR) CLNS Convergence Sublayer
Adaptation Layer
Segmentation & Reassembly Sublayer
MAC
802.6 segments/ ATM cells
Service Data Unit (SDU)
Initial MAC PDU (IMPDU)
48 octets.
unused
Derived MAC PDU (DMPDU)
53 octets

George H. Clapp

AMERITECH SERVICES

SMDS Protocol Layers

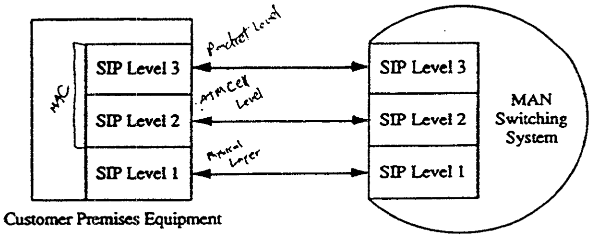

George H. Clapp

AMERITECH SERVICES

SMDS and 802.6 Protocol Layers

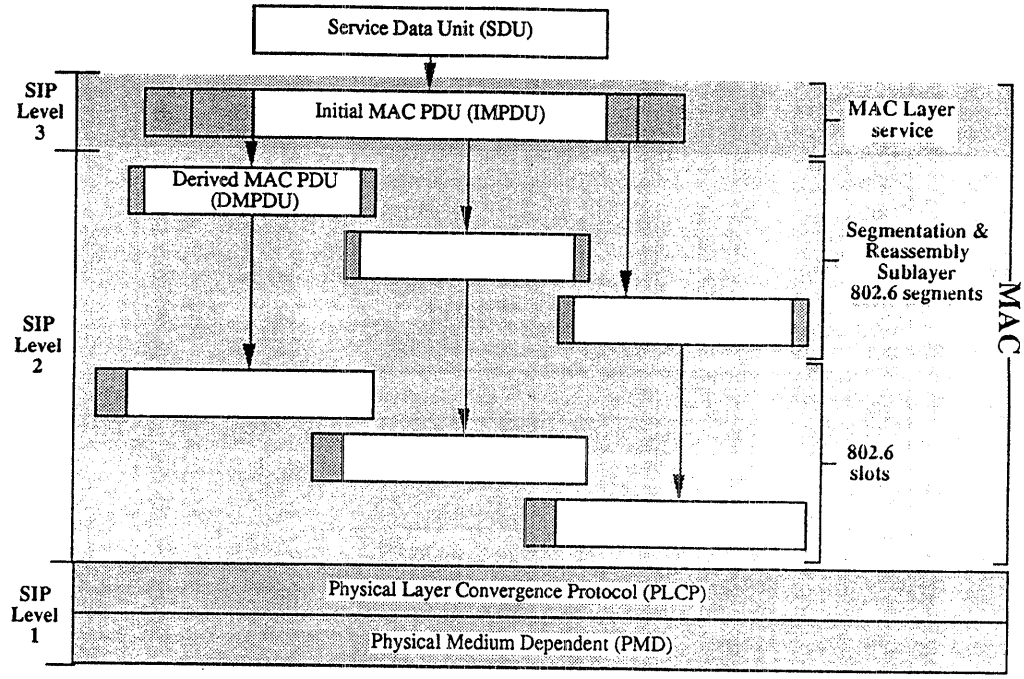

Service Data Unit (SDU)

SIP Level 3

Initial MAC PDU (IMPDU)

MAC Layer service

Derived MAC PDU (DMPDU)

Segmentation & Reassembly Sublayer 802.6 segments

MAC

SIP Level 2

802.6 slots

SIP Level 1

Physical Layer Convergence Protocol (PLCP)

Physical Medium Dependent (PMD)

802.6 and BISDN Protocol Layers

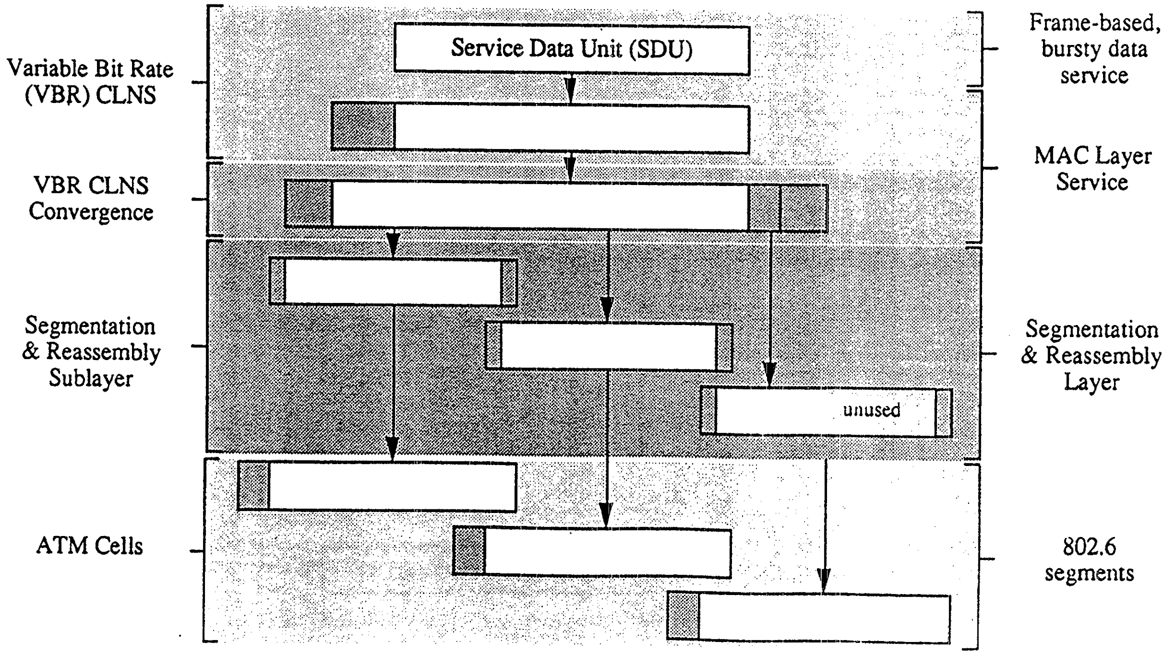

Variable Bit Rate (VBR) CLNS

Service Data Unit (SDU)

Frame-based, bursty data service

VBR CLNS Convergence

MAC Layer Service

Segmentation & Reassembly Sublayer

unused

Segmentation & Reassembly Layer

ATM Cells

802.6 segments

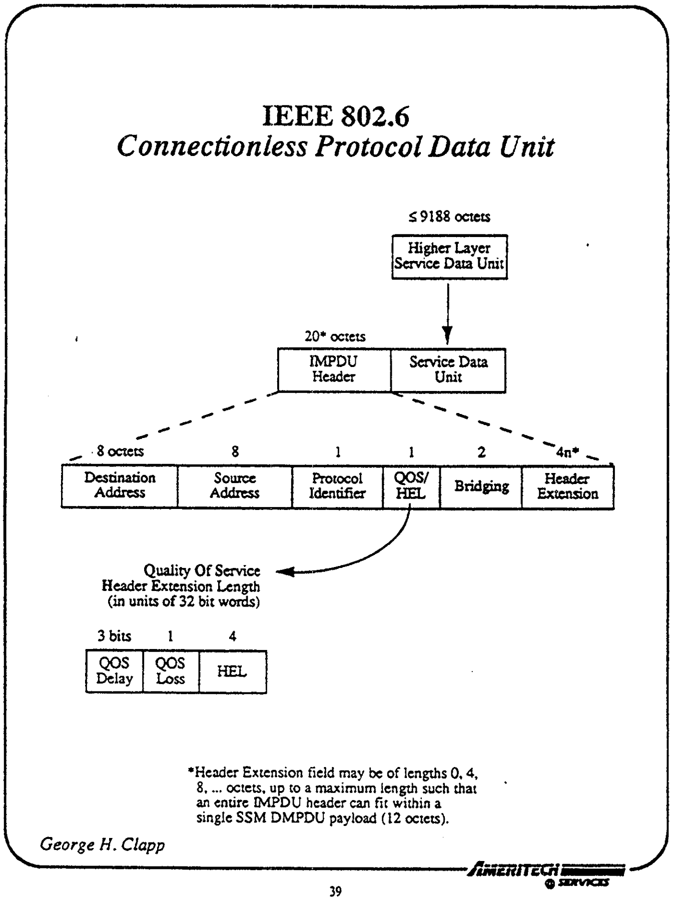

IEEE 802.6
Connectionless Protocol Data Unit

≤ 9188 octets

Higher Layer Service Data Unit

20* octets

IMPDU Header	Service Data Unit

Destination Address	Source Address	Protocol Identifier	QOS/ HEL	Bridging	Header Extension
8 octets	8	1	1	2	4n*

Quality Of Service Header Extension Length (in units of 32 bit words)

QOS Delay	QOS Loss	HEL
3 bits	1	4

*Header Extension field may be of lengths 0, 4, 8, ... octets, up to a maximum length such that an entire IMPDU header can fit within a single SSM DMPDU payload (12 octets).

George H. Clapp

AMERITECH SERVICES

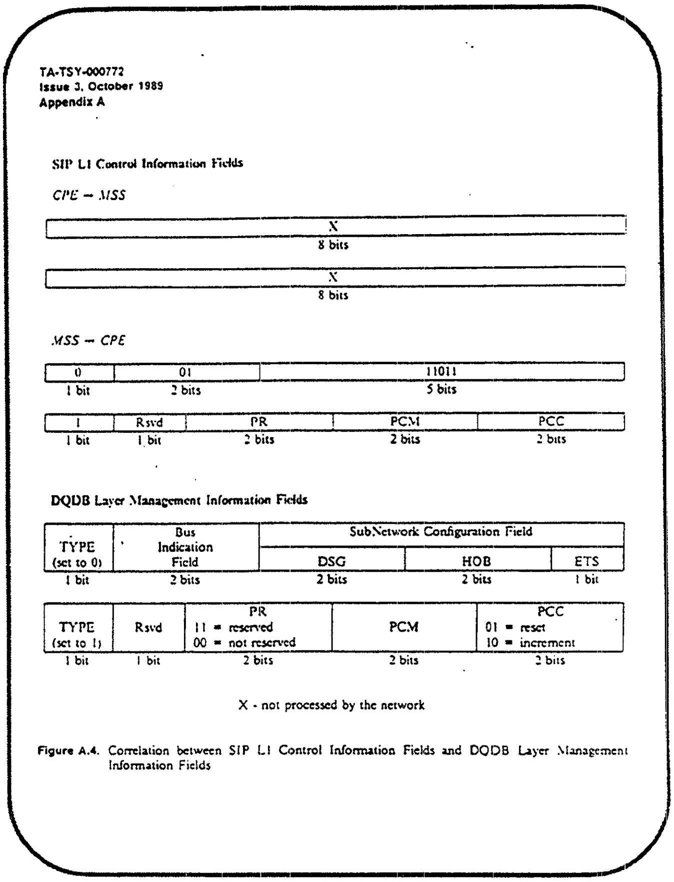

TA-TSY-000772
Issue 3, October 1989
Appendix A

SIP L1 Control Information Fields

CPE → MSS

X
8 bits

X
8 bits

MSS → CPE

0	01	11011
1 bit	2 bits	5 bits

1	Rsvd	PR	PCM	PCC
1 bit	1 bit	2 bits	2 bits	2 bits

DQDB Layer Management Information Fields

TYPE (set to 0)	Bus Indication Field	SubNetwork Configuration Field		
		DSG	HOB	ETS
1 bit	2 bits	2 bits	2 bits	1 bit

TYPE (set to 1)	Rsvd	PR 11 = reserved 00 = not reserved	PCM	PCC 01 = reset 10 = increment
1 bit	1 bit	2 bits	2 bits	2 bits

X - not processed by the network

Figure A.4. Correlation between SIP L1 Control Information Fields and DQDB Layer Management Information Fields

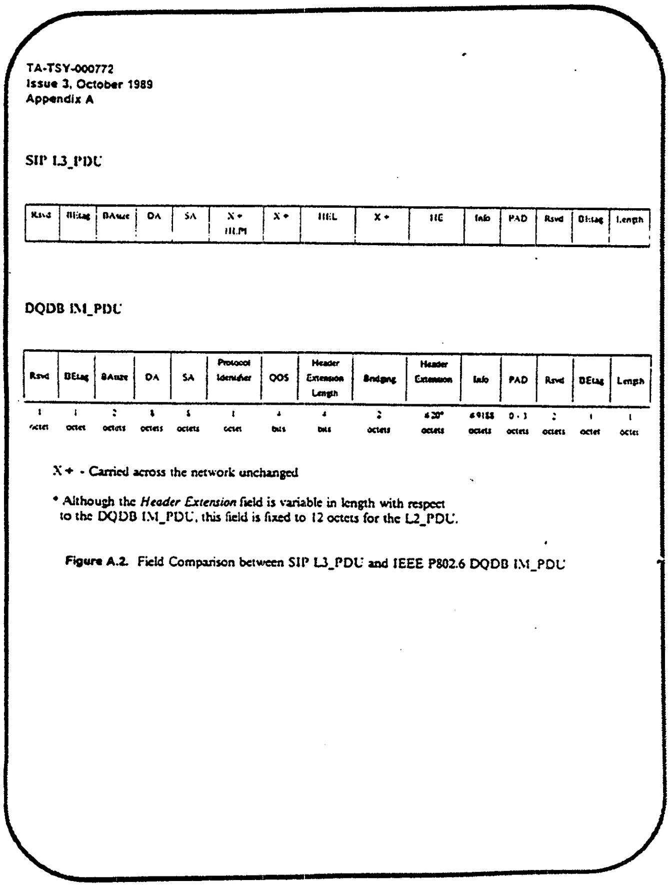

TA-TSY-000772
Issue 3, October 1989
Appendix A

SIP L3_PDU

Rsvd	BEtag	BAsize	DA	SA	X+ HLPI	X+	HEL	X+	HE	Info	PAD	Rsvd	BEtag	Length

DQDB IM_PDU

Rsvd	BEtag	BAsize	DA	SA	Protocol Identifier	QOS	Header Extension Length	Bridging	Header Extension	Info	PAD	Rsvd	BEtag	Length
1 octet	1 octet	2 octets	8 octets	8 octets	1 octet	4 bits	4 bits	2 octets	≤20* octets	≤9188 octets	0 - 3 octets	2 octets	1 octet	1 octet

X+ - Carried across the network unchanged

* Although the *Header Extension* field is variable in length with respect to the DQDB IM_PDU, this field is fixed to 12 octets for the L2_PDU.

Figure A.2. Field Comparison between SIP L3_PDU and IEEE P802.6 DQDB IM_PDU

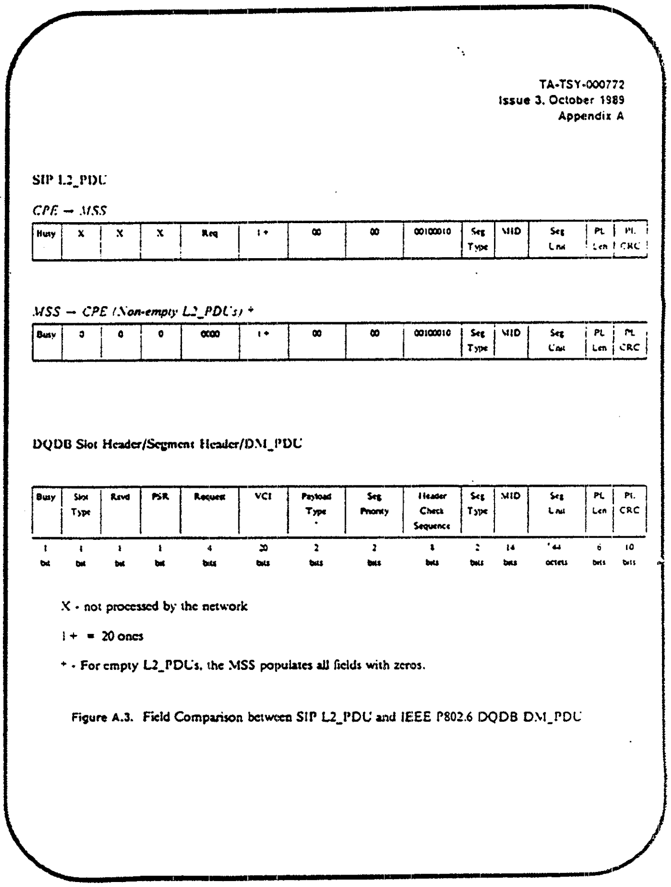

TA-TSY-000772
Issue 3, October 1989
Appendix A

SIP L2_PDU

CPE → MSS

Busy	X	X	X	Req	1+	00	00	00100010	Seg Type	MID	Seg Unit	PL Len	PL CRC

MSS → CPE (Non-empty L2_PDUs) +

Busy	0	0	0	0000	1+	00	00	00100010	Seg Type	MID	Seg Unit	PL Len	PL CRC

DQDB Slot Header/Segment Header/DM_PDU

Busy	Slot Type	Rsvd	PSR	Request	VCI	Payload Type	Seg Priority	Header Check Sequence	Seg Type	MID	Seg Unit	PL Len	PL CRC
1 bit	1 bit	1 bit	1 bit	4 bits	20 bits	2 bits	2 bits	8 bits	2 bits	14 bits	44 octets	6 bits	10 bits

X - not processed by the network

1+ = 20 ones

+ - For empty L2_PDUs, the MSS populates all fields with zeros.

Figure A.3. Field Comparison between SIP L2_PDU and IEEE P802.6 DQDB DM_PDU

IEEE 802.6
Address Formats

4 bits	60-N bits	N bits
Address_Type	Padding	Address

Address_Type

0100	16 bit IEEE 802 Address
1000	48 bit IEEE 802 Address
1100	60 bit, Individual, Publicly Administered
1101	60 bit, Individual, Privately Administered
1110	60 bit, Group, Publicly Administered
1111	60 bit, Group, Privately Administered

CCITT E.164 address (ISDN telephone number)

north american

Country Code	National Destination Code	Subscriber Number
0001		

Use of E.164 address facilitates interworking with the public network.

Variable length, up to 15 digits (0-9) encoded in Binary Coded Decimal.
Country Code: 1-3 digits
National Destination & Subscriber Number vary in length depending on requirements of the destination country.

IEEE 802 48 bit address

Individual/ Group	Universally/ Locally Administered	46 bit address

IEEE 802 16 bit address

Individual/ Group	15 bit address

George H. Clapp

Adaptation Layer

Layer of functionality which "adapts" a non-cell-based service (e.g., 802.6 connectionless packet service) to the cell-based ATM transport.

IEEE 802.6 Working Group and CCITT have accepted a common adaptation layer for the "Variable Bit Rate (VBR) Connectionless Network Service (CLNS)."

Two logical sublayers:

Variable Bit Rate (VBR) Connectionless Network Service (CLNS) Convergence Protocol Sublayer.

Segmentation and Reassembly (SAR) Sublayer with error control.

George H. Clapp

Adaptation Layer
VBR CLNS Convergence Protocol Sublayer

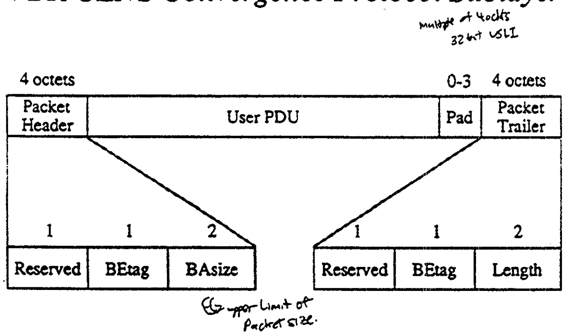

BEtag Same value is placed in the header and trailer fields; used to associate header and trailer of the same PDU for error control.

BAsize Used by receiver for buffer management; either...

Length in octets of 802.6 IMPDU (header and information, inclusive), or...

Greater than or equal to the true PDU length.

Length Length in octets of the user PDU (less the Pad).

Pad A 0 to 3 octet field added to the end of the user PDU to align the Packet Trailer to a 32-bit boundary.

George H. Clapp

Adaptation Layer
Segmentation and Reassembly (SAR) Sublayer

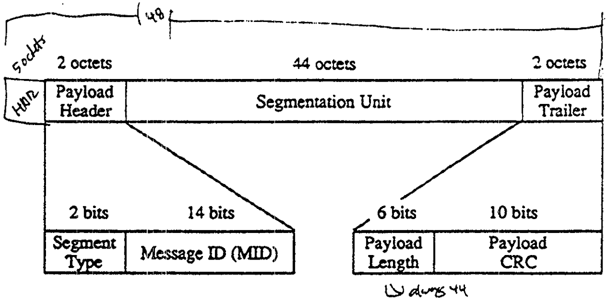

Segment Type Used for packet delineation;
Encoding

Beginning of Message (BOM):	10
Continuation of Message (COM):	00
End of Message (EOM):	01
Single Segment Message (SSM):	11

MID Used to reassemble segments into packets; all cells of a given packet will have the same MID value.

Payload Length Number of octets of packets included in the payload of the segmentation unit (1-44) (4-44 for 802.6 CL MAC service).

Payload CRC CRC calculation over the entire contents of the segment payload, including payload header and *payload length*. Error detection is mandatory and single bit error correction is optional.

Generating polynomial: $G(x) = x^{10} + x^{9} + x^{5} + x^{4} + x + 1$ R.

George H. Clapp

Adaptation Layer
CRC per Cell

IEEE 802 performance objective: MAC layer shall not deliver more than one errored frame per year.

error control in ATM cell

CRC per Cell...

- provides sufficient error control to meet IEEE 802 performance specifications;
- eliminates the need for error control at higher layers;
- supports simple, "light weight", fast transport protocols.

Destination Address	Source Address	Protocol Identifier	QOS/ HEL	Bridging	Header Extension	Info

802.6 Connectionless IMPDU format.

CRC per Cell
Sources of Error

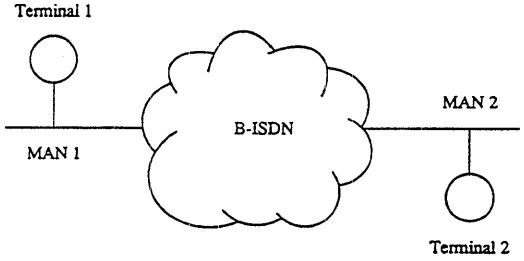

3 sources of error from terminal 1 to terminal 2 in a fiber optic network:†

Random Errors: 10^{-12}.

Burst errors (protection switching): 0.24 events per day on 1000 mile system; 20-40 ms duration.

Buffer overflow: engineering parameter.

† DS-1 Facility Performance: Optical Fiber as a Long Haul Technology, K. A. Tse, R. M. O'Connor & N. Fatseas, ICC '87, pp. 646-649.

CRC per Cell
Undetected Errors

Components of undetected errors.

1. Errors in all "errored cells" are not detected (Dominant).
2. Last cell in error, discovered to be in error, dropped, cells in the next frame are errored in such a way that the collected cells appear to be a legitimate frame.

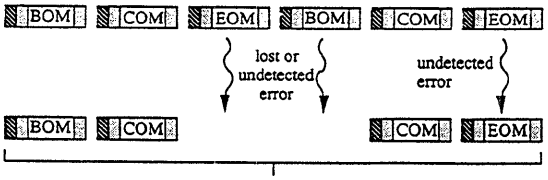

MID's match, BEtag's match, & LENGTH is correct.

3. At least one COM cell is lost and LENGTH field of EOM cell is errored to match truncated frame.

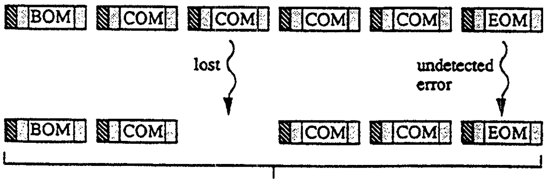

Errored LENGTH is correct.

CRC per Cell
Comparison vs. Bit Error Rate

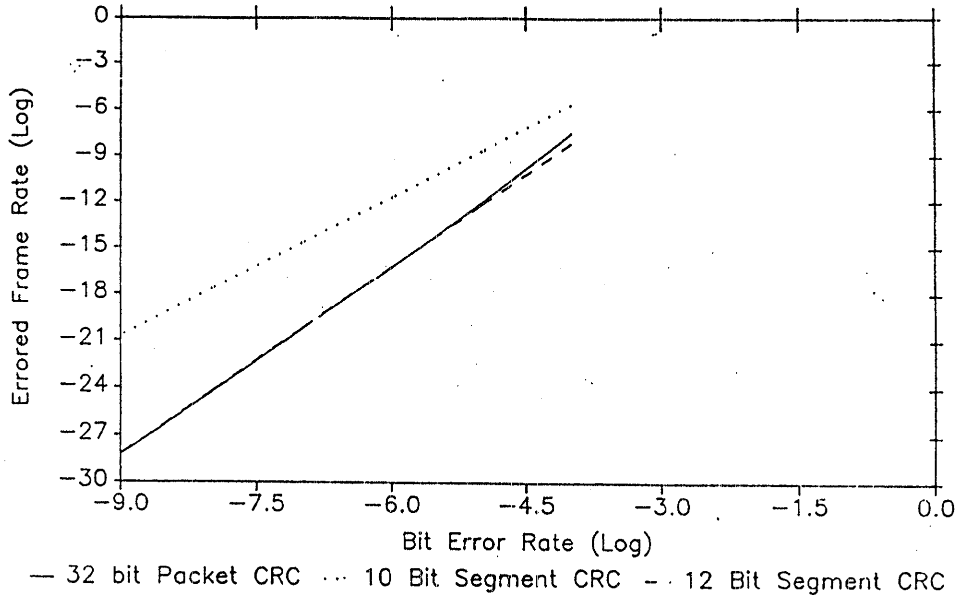

A Packet Size of 100 Segments is Assumed

George H. Clapp

CRC per Cell
Comparison vs. Packet Length

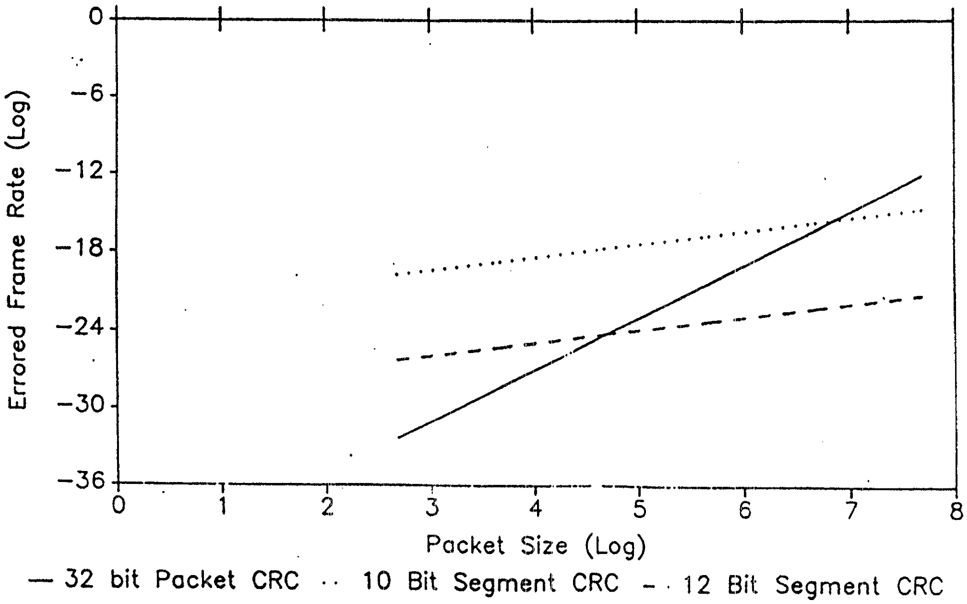

A Bit Error Rate of 10E-8 is Assumed

George H. Clapp

ATM Cell Format

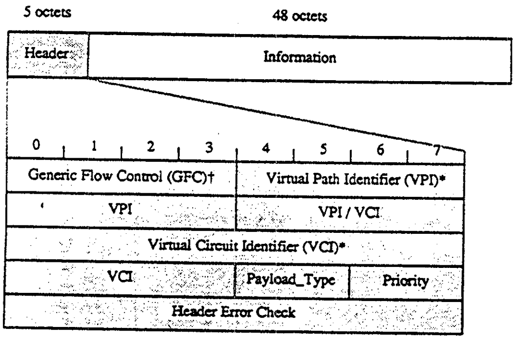

† At the Network-Network Interface (NNI), the GFC does not exist; the VPI/VCI fields extend into this field.

* At User-Network Interface (UNI), VPI and VCI are either 8 or 12 bits; no more than 20 bits of the combined VPI/VCI fields are significant at any time.

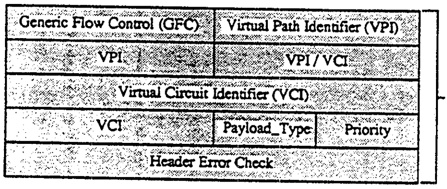
Generic Flow Control (GFC)
Virtual Path Identifier (VPI)
VPI
VPI / VCI
Virtual Circuit Identifier (VCI)
VCI
Payload_Type
Priority
Header Error Check

HCS coverage;
Generating polynomial
x^8+x^2+x+1

George H. Clapp

802.6 Segment Format

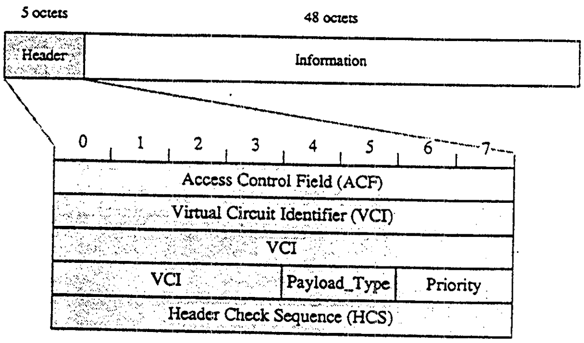
5 octets
48 octets
Header
Information
0
1
2
3
4
5
6
7
Access Control Field (ACF)
Virtual Circuit Identifier (VCI)
VCI
VCI
Payload_Type
Priority
Header Check Sequence (HCS)

by not covering the ACF, reduce latency ⇒ pipelining possible.

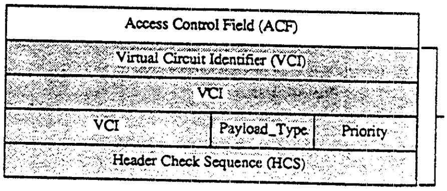
Access Control Field (ACF)
Virtual Circuit Identifier (VCI)
VCI
VCI
Payload_Type
Priority
Header Check Sequence (HCS)

HCS coverage;
Generating polynomial
x^8+x^2+x+1

George H. Clapp

IEEE 802.6
Access Control Field

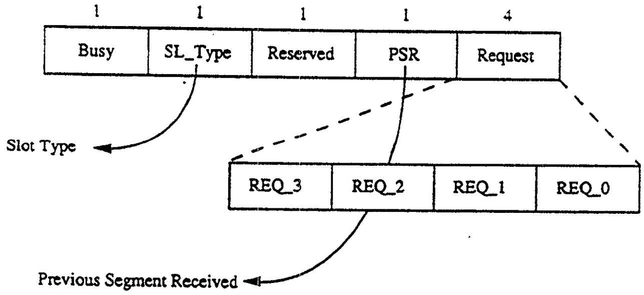

Busy	SL_Type	Slot State
0	0	Empty, Queue Arbitrated slot
0	1	Invalid
1	0	Busy, Queue Arbitrated slot
1	1	Pre-Arbitrated slot

George H. Clapp

Virtual Path Concept

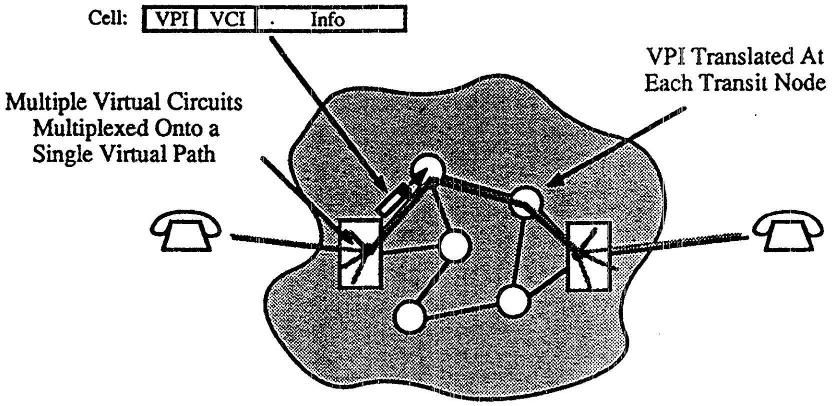

Cell: VPI | VCI | . Info

VPI Translated At Each Transit Node

Multiple Virtual Circuits Multiplexed Onto a Single Virtual Path

Advantages:

- Removes Necessity of Per-Call Processing and Associated Data from Transit Exchanges.
- Faster Call Establishment.
- More Flexible Network Architectures.

Header Check Sequence

Access Control Field (ACF) and Virtual Path ...	
Identifier (VPI)	Virtual Circuit Identifier
VCI	
VCI	Payload_Type · Priority
HCS	

HCS coverage; Generating polynomial x^8+x^2+x+1

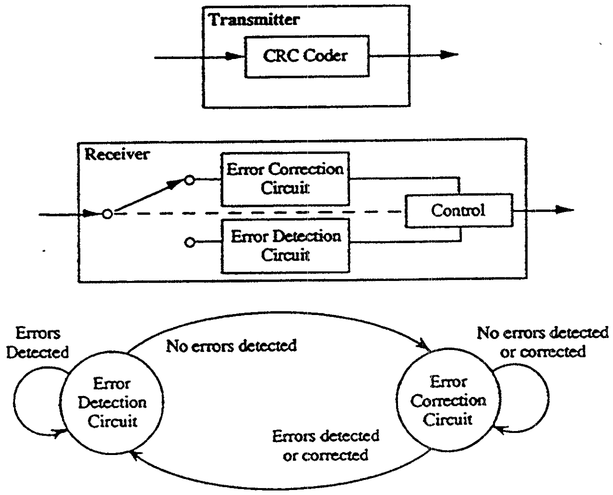

Transmitter

CRC Coder

Receiver

Error Correction Circuit

Control

Error Detection Circuit

Errors Detected

No errors detected

No errors detected or corrected

Error Detection Circuit

Error Correction Circuit

Errors detected or corrected

Transitions occur upon reading a segment.

IEEE 802.6 Protocol Structures
"Quick Reference Card 1"

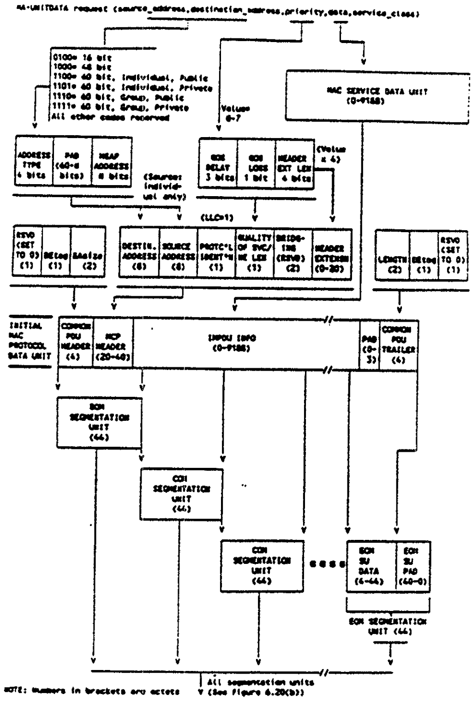
MA-UNITDATA request (source_address,destination_address,priority,data,service_class)
0100= 16 bit
1000= 48 bit
1100= 60 bit, Individual, Public
1101= 60 bit, Individual, Private
1110= 60 bit, Group, Public
1111= 60 bit, Group, Private
All other codes reserved
Value 0-7
MAC SERVICE DATA UNIT (0-9188)
ADDRESS TYPE 4 bits
PAD (60-n bits)
MEAP ADDRESS n bits
(Source: individual only)
QOS DELAY 3 bits
QOS LOSS 1 bit
HEADER EXT LEN 4 bits
(Value x 4)
(LLC=1)
RSVD (SET TO 0) (1)
BEtag (1)
BAsize (2)
DESTIN. ADDRESS (8)
SOURCE ADDRESS (8)
PROTC'L IDENT'N (1)
QUALITY OF SVC/ HE LEN (1)
BRIDGING (RSVD) (2)
HEADER EXTENSN (0-20)
LENGTH (2)
BEtag (1)
RSVD (SET TO 0) (1)
INITIAL MAC PROTOCOL DATA UNIT
COMMON PDU HEADER (4)
MCP HEADER (20-40)
IMPDU INFO (0-9188)
PAD (0-3)
COMMON PDU TRAILER (4)
BOM SEGMENTATION UNIT (44)
COM SEGMENTATION UNIT (44)
COM SEGMENTATION UNIT (44)
EOM SU DATA (4-44)
EOM SU PAD (40-0)
EOM SEGMENTATION UNIT (44)
All segmentation units (See Figure 6.20(b))
NOTE: Numbers in brackets are octets

IEEE 802.6 Protocol Structures
"Quick Reference Card 2"

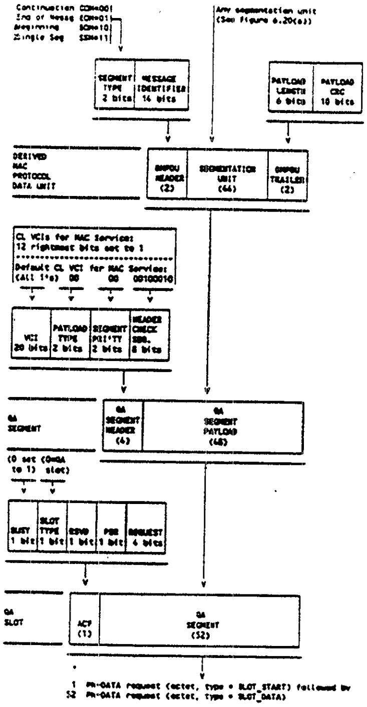

Frame Reassembly

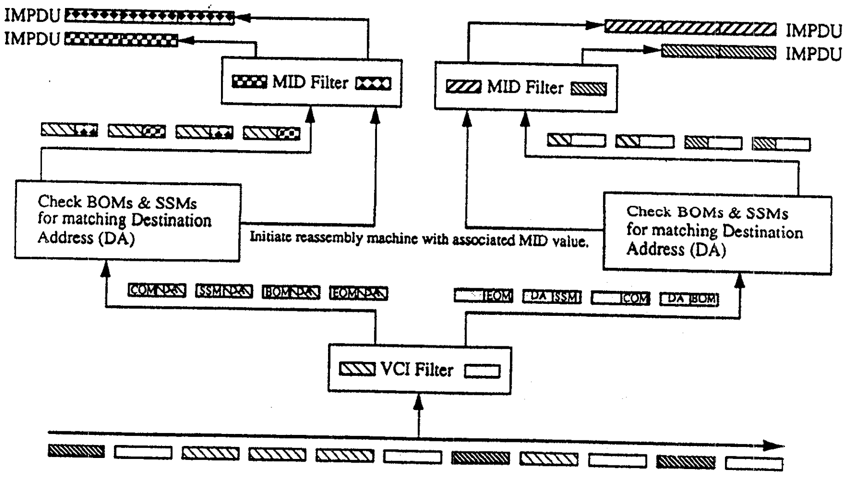
IMPDU
IMPDU
MID Filter
IMPDU
IMPDU
MID Filter
Check BOMs & SSMs for matching Destination Address (DA)
Initiate reassembly machine with associated MID value.
Check BOMs & SSMs for matching Destination Address (DA)
COM/DA
SSM/DA
BOM/DA
EOM/DA
EOM
DA SSM
COM
DA BOM
VCI Filter

Frame Reassembly
Flowchart

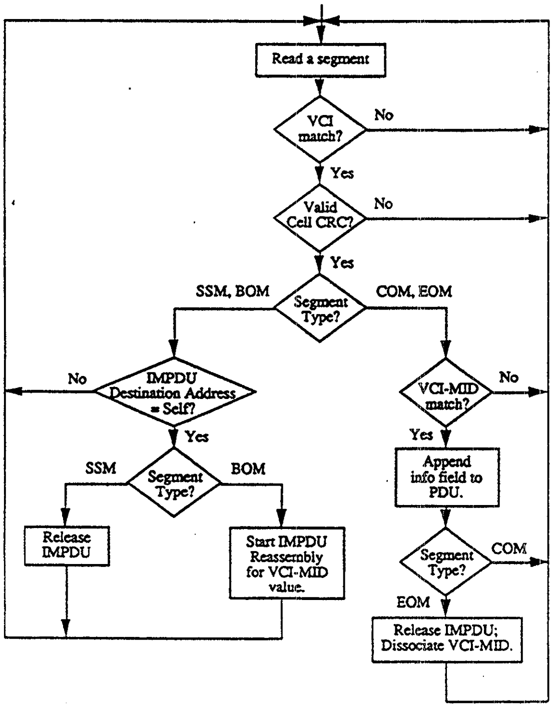

Read a segment

VCI match?

No

Yes

Valid Cell CRC?

No

Yes

SSM, BOM

Segment Type?

COM, EOM

IMPDU Destination Address = Self?

No

Yes

SSM

Segment Type?

BOM

Release IMPDU

Start IMPDU Reassembly for VCI-MID value.

VCI-MID match?

No

Yes

Append info field to PDU.

Segment Type?

COM

EOM

Release IMPDU; Dissociate VCI-MID.

Destination Release

Without Destination Release, segments used by Station 1 to send to Station 2 are not reused further downstream.

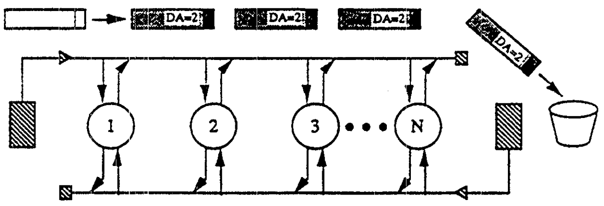

With Destination Release, segments can be reused, increasing network throughput.

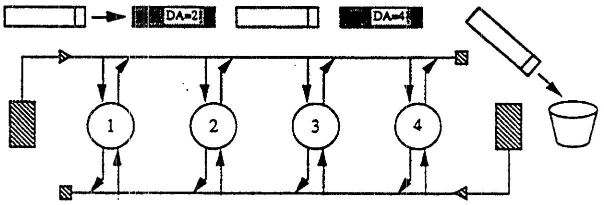

Destination Release
Previous Segment Received Bit

To recognize whether a segment is destined to itself, a station must examine the Destination Address within the IMPDU header.

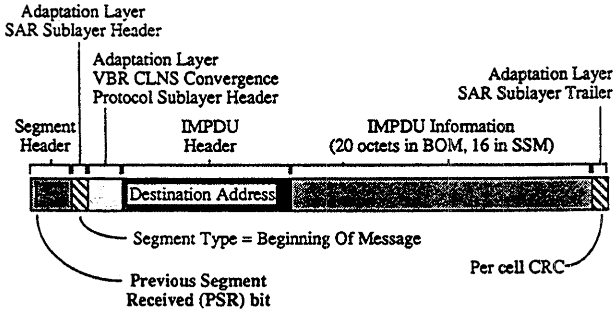

Adaptation Layer SAR Sublayer Header

Adaptation Layer VBR CLNS Convergence Protocol Sublayer Header

Adaptation Layer SAR Sublayer Trailer

Segment Header

IMPDU Header

IMPDU Information (20 octets in BOM, 16 in SSM)

Destination Address

Segment Type = Beginning Of Message

Per cell CRC

Previous Segment Received (PSR) bit

Destination Address ends 19 octets inside the segment.

Incur excessive station latency to set a bit within the current segment, so the PSR bit in the *following* segment is set.

Destination Release
Erasure Node

Special "Erasure Node" buffers an entire segment plus the ACF of the following segment. If the PSR bit in the following segment is set, the "previous" segment is "erased" and transmitted for reuse further downstream.

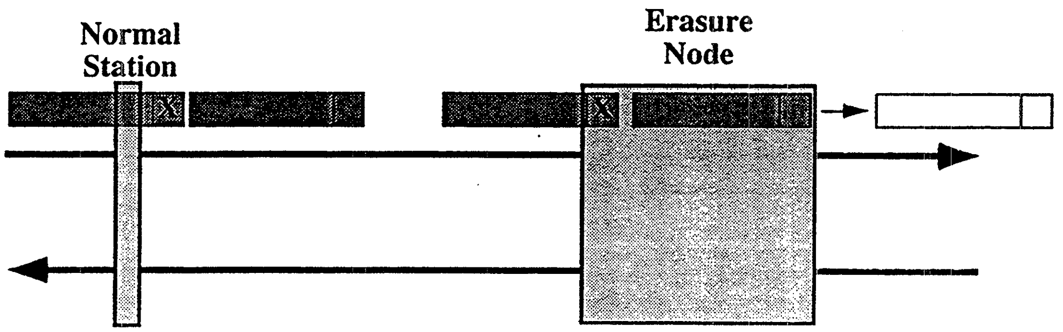

DQDB
Message ID Allocation

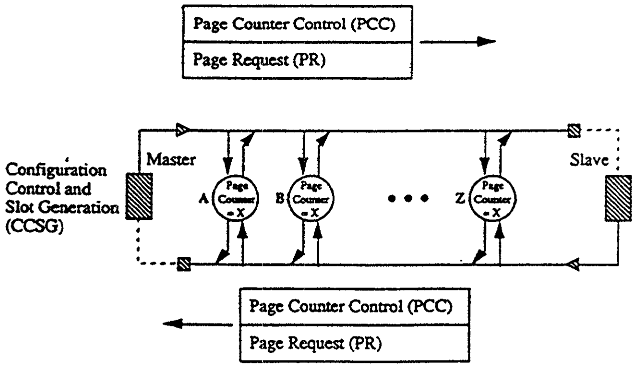

- MID address space ranges between MIN_PAGE and MAX_PAGE. Each page contains PAGE_SIZE (4) MID values (one per REQ priority).
- Each station maintains a "Page Counter" and cycles through the address space under control of the PCC parameter issued by Master CCSG.

  PCC = INCREMENT, station adds one to Page Counter.

  PCC = RESET, station sets Page Counter to MIN_PAGE.
- A Page Request (PR) parameter is associated with each page value .
- Slave CCSG transparently "wraps around" PCC and PR parameters.

DQDB MID Allocation
Two-Pass Algorithm
"Keep" Pass

Station B "owns" Page 10...

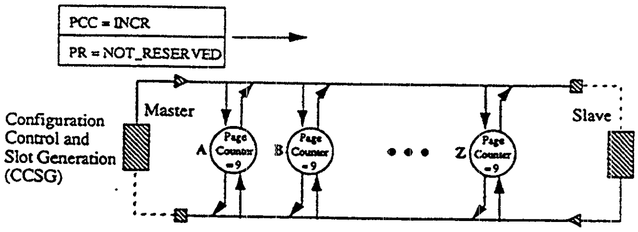

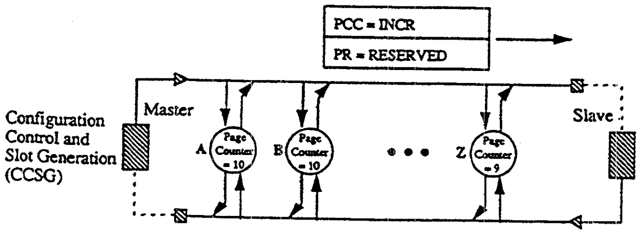

DQDB MID Allocation

Two-Pass Algorithm
"Get" Pass

Station B "wants" Page 15...

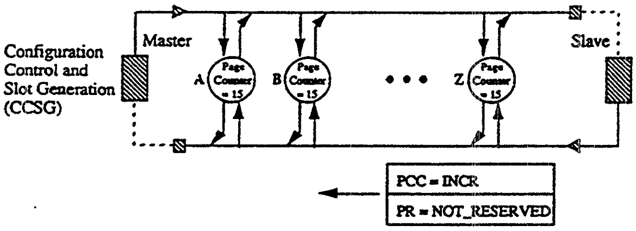

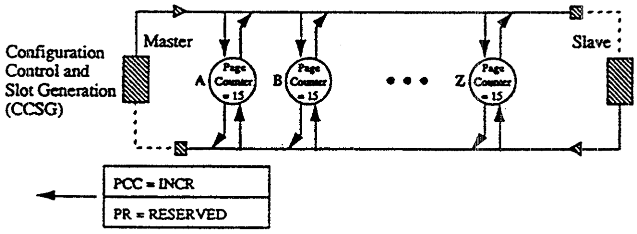

DQDB Performance

The Parameter "a"

Parameter "a" is a measure of the number of bits "in flight" in a network. Factors are...

- Length of the network and propagation delay of the medium,
- Transmission speed,
- Station latency.

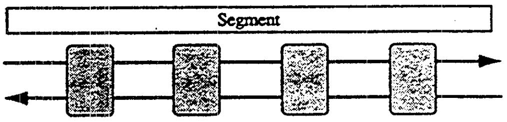

"a" = 1 (measured in slots)

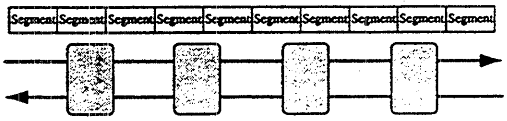

"a" = 10

DQDB Performance

Average Access Delay

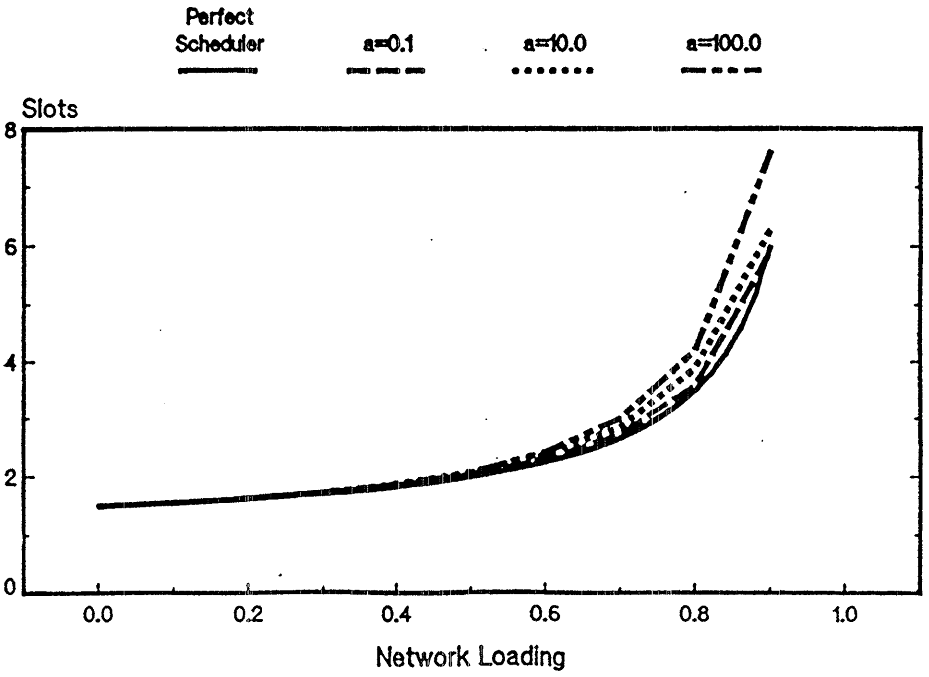

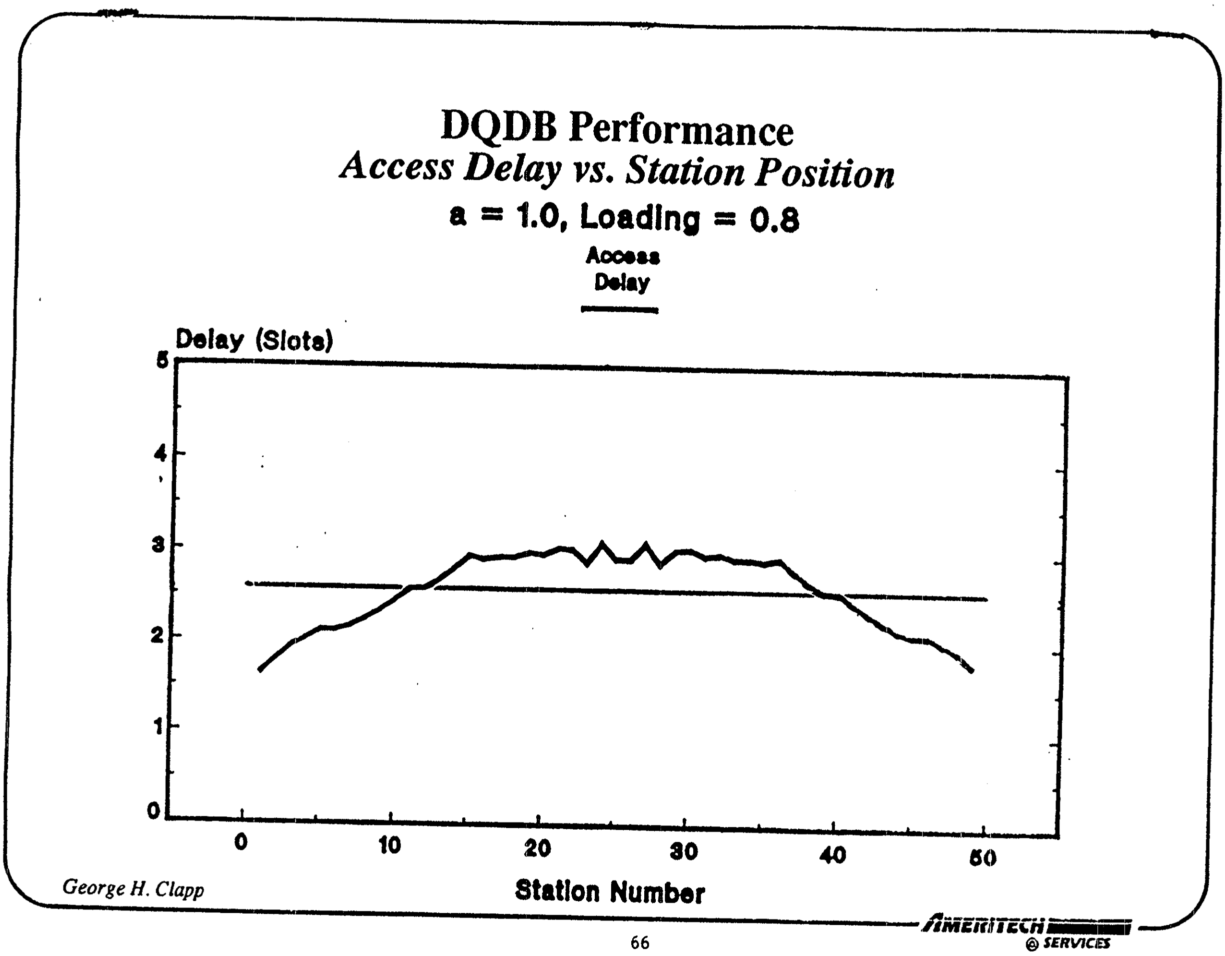

DQDB Performance

Access Delay vs. Station Position

a = 1.0, Loading = 0.8

George H. Clapp

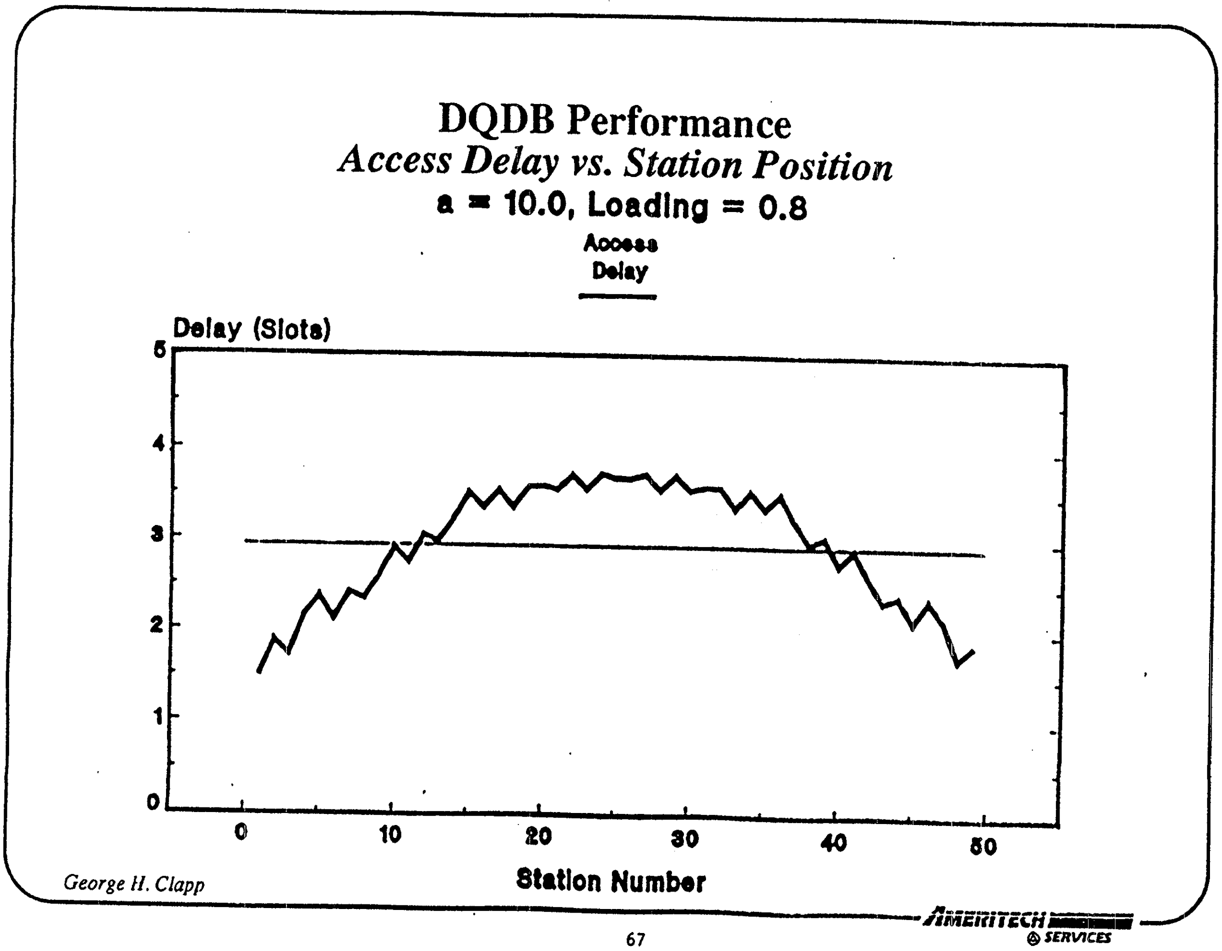

DQDB Performance

Access Delay vs. Station Position

a = 10.0, Loading = 0.8

George H. Clapp

DQDB Performance

Access Delay vs. Station Position

a = 100.0, Loading = 0.1

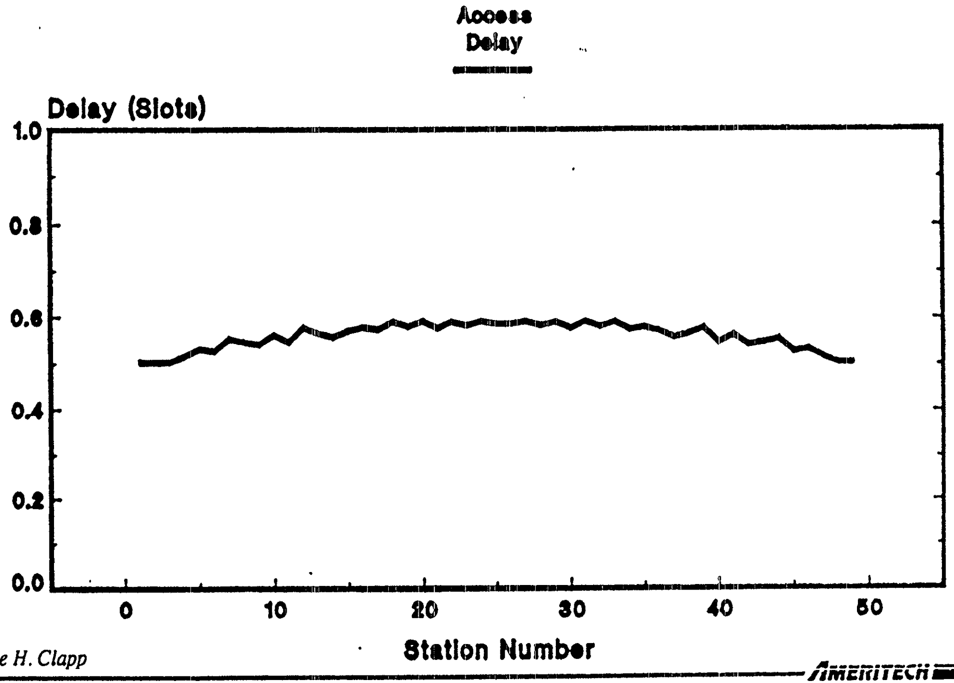

George H. Clapp

DQDB Performance

Access Delay vs. Station Position

a = 100.0, Loading = 0.5

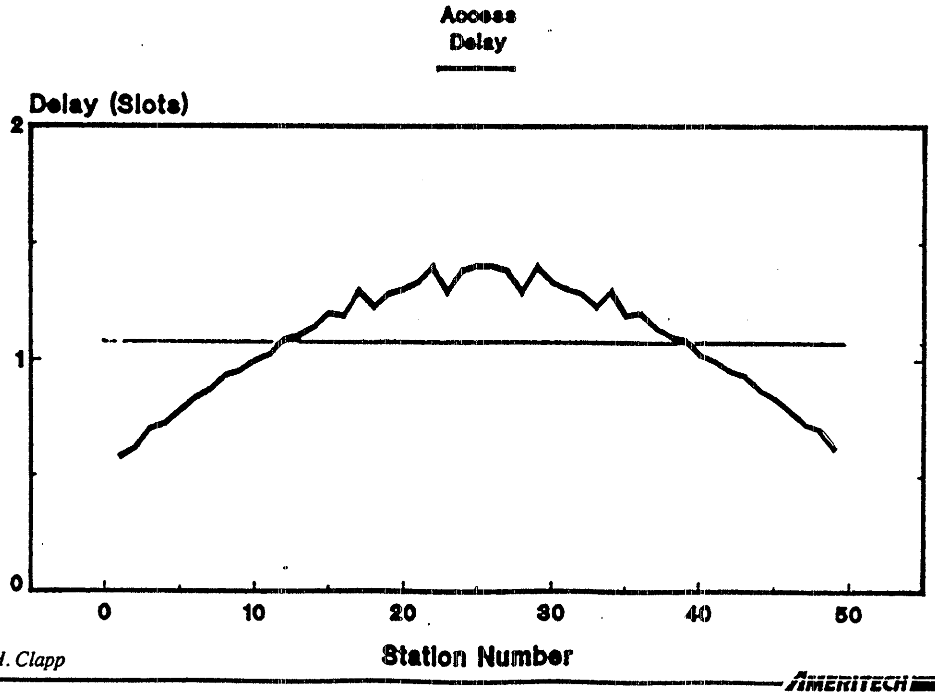

George H. Clapp

DQDB Performance

Access Delay vs. Station Position

a = 100.0, Loading = 0.8

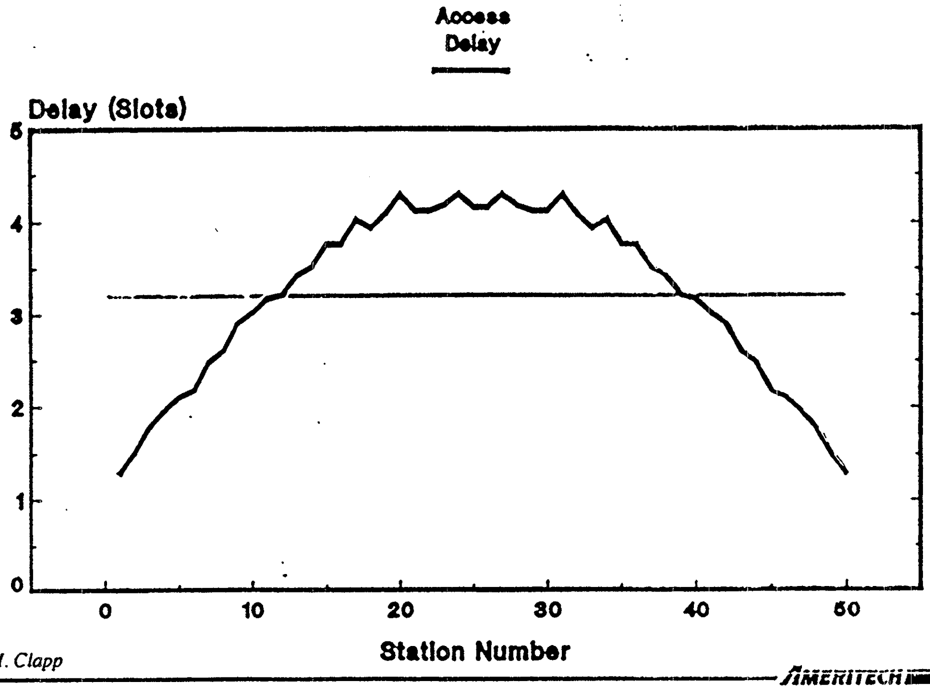

George H. Clapp

Physical Layer

Primitives

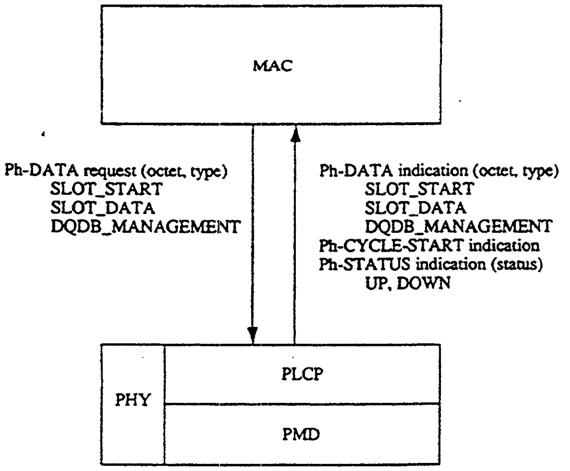

Ph-DATA request (octet, type)
SLOT_START
SLOT_DATA
DQDB_MANAGEMENT

Ph-DATA indication (octet, type)
SLOT_START
SLOT_DATA
DQDB_MANAGEMENT
Ph-CYCLE-START indication
Ph-STATUS indication (status)
UP, DOWN

George H. Clapp

Physical Layer Convergence Protocol (PLCP)

Higher Layers	
Data Link Layer	Logical Link Control (LLC)
	Medium Access Control (MAC)
Physical (PHY)	**Physical Layer Convergence Protocol (PLCP)**
	Physical Medium Dependent (PMD)

Functions of PLCP:

- Allows DQDB Layer to operate independently of the nature of the transmission system.
- "Maps" segments/cells/DMPDUs onto the payload of the transmission system.
- Derives 125 µsec signal (Ph-CYCLE-START) to the MAC from the PMD.
- Transport of network management information and of transmission system Operations, Administration, and Management (OA&M) information.

George H. Clapp

Digital Hierarchy Bit Rates
(CCITT G.702)

Digital Hierarchy Level	North American Bit Rates	Japanese Bit Rates	European (CEPT†) Bit Rates
0	64 Kb/s	64 Kb/s	64 Kb/s
1	1.544 Mb/s	1.544 Mb/s	2.048 Mb/s
2	6.312 Mb/s	6.312 Mb/s	8.448 Mb/s
3	44.736 Mb/s	32.064 Mb/s	34.368 Mb/s
4		97.728 Mb/s	139.264 Mb/s

†European Conference of Post and Telecommunication Administrations

George H. Clapp

North American Digital Signal (DS) Transmission Hierarchy

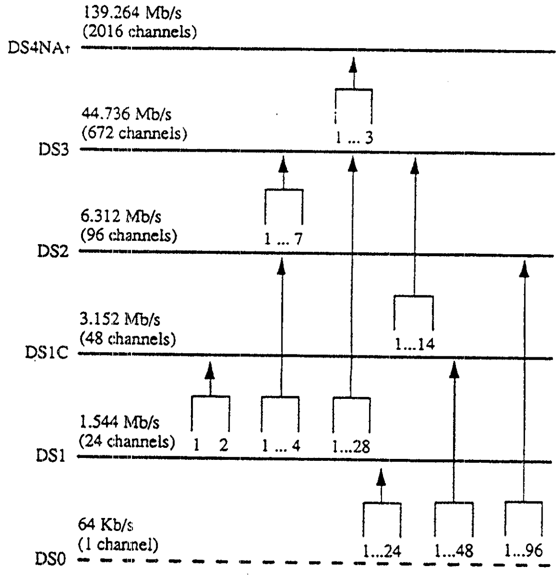

DS4NA† 139.264 Mb/s (2016 channels)

DS3 44.736 Mb/s (672 channels)

DS2 6.312 Mb/s (96 channels)

DS1C 3.152 Mb/s (48 channels)

DS1 1.544 Mb/s (24 channels)

DS0 64 Kb/s (1 channel)

George H. Clapp

† DS4NA: Digital Signal 4 North American

DS1 Transmission
Extended Superframe

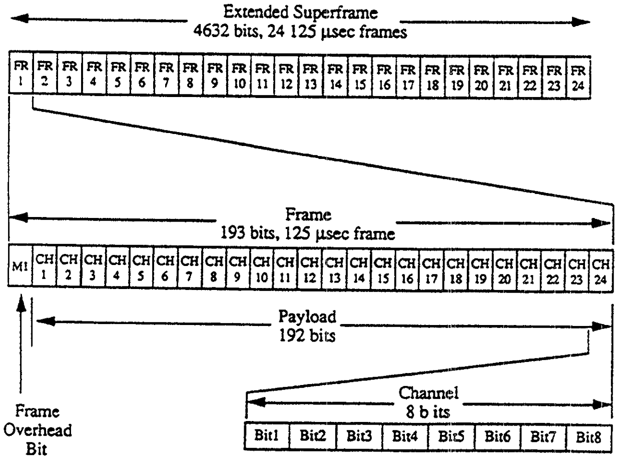

Extended Superframe
4632 bits, 24 125 μsec frames

Frame
193 bits, 125 μsec frame

Payload
192 bits

Channel
8 b its

Frame Overhead Bit

1.536 Mb/s payload.

First bit of each frame is used for transmission overhead.

George H. Clapp

Proposed DS1 Mapping
SMDS

3 Superframe's Per Frame

1.17 mb/s ⇒ for user segment.

A1	A2	S1	Z1	802.6 segment 1	
A1	A2	S2	Z2	802.6 segment 2	
A1	A2	S3	B1	802.6 segment 3	
A1	A2	S4	G1	802.6 segment 4	
A1	A2	S5	F1	802.6 segment 5	
A1	A2	S6	M1	802.6 segment 6	
A1	A2	S7	M2	802.6 segment 7	
A1	A2	S8	T1	802.6 segment 8	
A1	A2	S9	T2	802.6 segment 9	
A1	A2	S10	C1	802.6 segment 10	6 octets

3 msec

A1	Framing Byte (F6 hex)	**F1**	PLCP Path User Channel
A2	Framing Byte (28 hex)	**Z1-Z2**	Growth
S1-S10	Segment counter	**M1-M2**	DQDB network management
B1	Bit Interleaved Parity 8 (BIP-8)	**T1-T2**	Timing Source Counter
G1	PLCP Path Status	**C1**	Cycle/Stuff Counter

George H. Clapp

AMERITECH SERVICES

DS3 Transmission

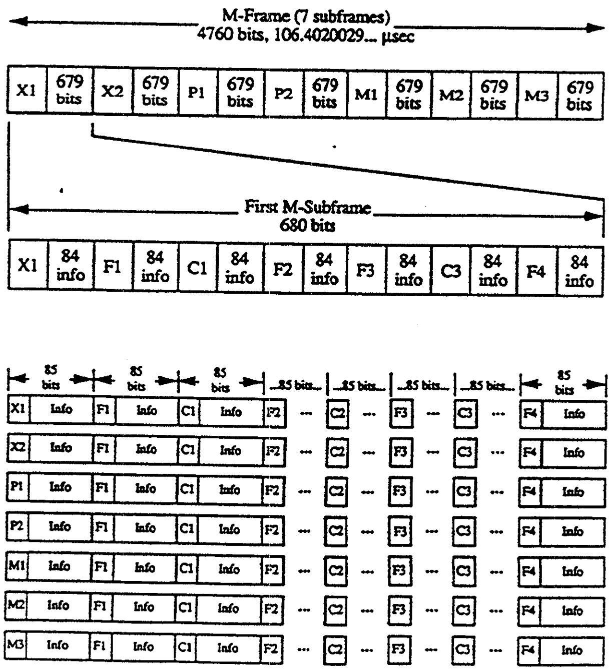

George H. Clapp

AMERITECH SERVICES

DS3 Mapping
IEEE 802.6 Working Group

Impose 125 µsec frame upon DS3 payload:

84/85 × 44.736 Mb/s	= 44.2097 Mb/s payload
44.2097 Mb/s × 125 µsec	= 5526.211765... bits
	= 690.776... octets
690.776 / 53 octets/cell	= 13 cells/cycle + 1.776 octet (too little)
12 slots × 48 × 8 / 125 µsec	= 36.864 Mb/s

4				53	
A1	A2	P11	Z6	DQDB Slot #1	
A1	A2	P10	Z5	DQDB Slot #2	
A1	A2	P9	Z4	DQDB Slot #3	
A1	A2	P8	Z3	DQDB Slot #4	
A1	A2	P7	Z2	DQDB Slot #5	
A1	A2	P6	Z1	DQDB Slot #6	
A1	A2	P5	F1	DQDB Slot #7	
A1	A2	P4	B1	DQDB Slot #8	
A1	A2	P3	G1	DQDB Slot #9	
A1	A2	P2	M2	DQDB Slot #10	
A1	A2	P1	M1	DQDB Slot #11	
A1	A2	P0	C1	DQDB Slot #12	13-14 nibbles

125 µsec

A1	Framing Octet (F6 hex)	**F1**	PLCP Path User Channel
A2	Framing Octet (28 hex)	**Z6-Z1**	Growth
P11-P0	Path Overhead Identifier Octets	**M2-M1**	Management Information
B1	Bit Interleaved Parity, BIP-8	**C1**	Cycle/Stuff Counter (375 µsec stuffing opportunity cycle)
G1	PLCP Path Status		

George H. Clapp

AMERITECH SERVICES

Synchronous Optical Network
SONET

Standardized rates and formats for an optical interface.

Synchronous optical hierarchy capable of carrying many different capacity signals.

Layered approach to transmission of OA&M information.

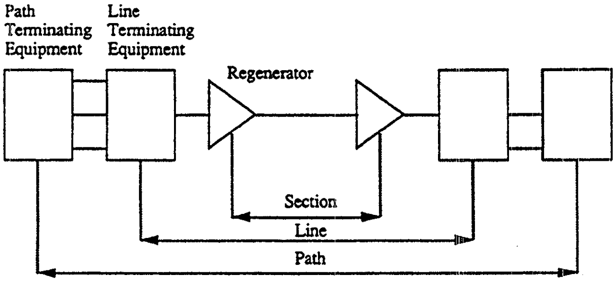

Section Repeater, regenerator.

Line Line terminating equipment, e.g., high-speed cross-connect, add/drop multiplexer, multiplexer.

Path Terminates SONET payload, e.g., DS1 cross-connect, DS1 add/drop multiplexer, DS1 multiplexer.

SONET
Synchronous Transport Signal level 1

STS-1: basic logical building block signal

OC-N: Optical Carrier level N (N×STS-1)

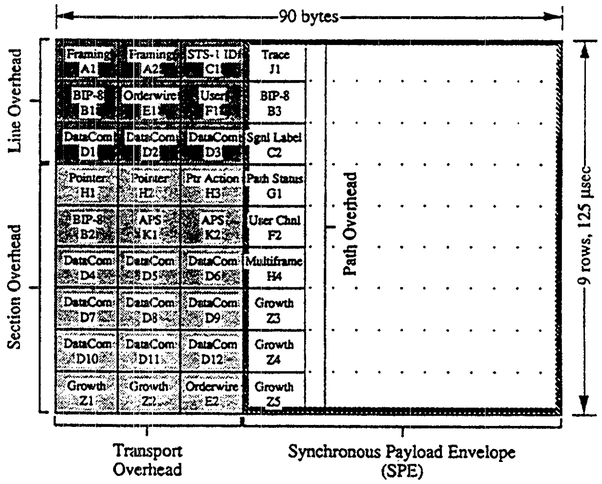

Framing A1	Framing A2	STS-1 ID C1	Trace J1
BIP-8 B1	Orderwire E1	User F1	BIP-8 B3
DataCom D1	DataCom D2	DataCom D3	Sgnl Label C2
Pointer H1	Pointer H2	Ptr Action H3	Path Status G1
BIP-8 B2	APS K1	APS K2	User Chnl F2
DataCom D4	DataCom D5	DataCom D6	Multiframe H4
DataCom D7	DataCom D8	DataCom D9	Growth Z3
DataCom D10	DataCom D11	DataCom D12	Growth Z4
Growth Z1	Growth Z2	Orderwire E2	Growth Z5

9 rows × 90 bytes × 8 / 125 µsec = 51.840 Mb/s line signaling rate.

9 rows × 87 bytes × 8 / 125 µsec = 50.112 Mb/s SPE rate.

9 rows × 86 bytes × 8 / 125 µsec = 49.536 Mb/s user data rate.

SONET
Plesiochronous Operation

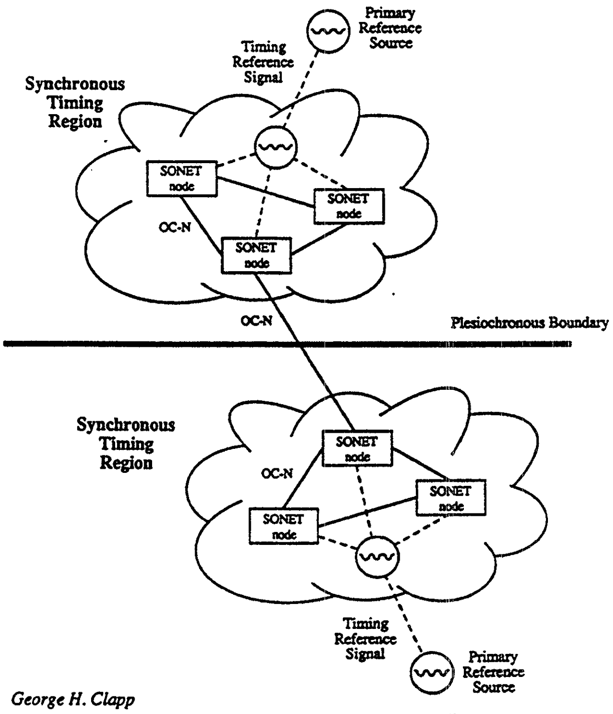

SONET
STS-1 Synchronous Payload Envelope Spanning STS-1 Frame

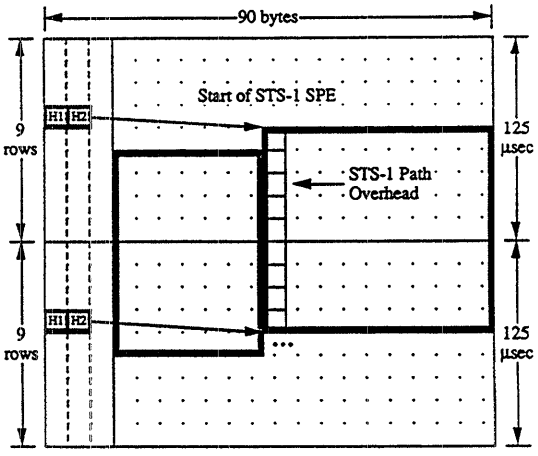

H1 & H2 pointers in Line Overhead used to adjust for "slips" and "stuffs" across plesiochronous boundaries.

SPE may begin anywhere within 125 µsec STS-1 frame.

SONET
Synchronous Hierarchical Rates

Multiplex N STS-1 signals into an STS-N signal.

OC Level	Line Rate (Mb/s)	OC Level	Line Rate (Mb/s)
OC-1	51.840	OC-18	933.120
OC-3	155.520	OC-24	1244.160
OC-9	466.560	OC-36	1866.240
OC-12	622.080	OC-48	2488.320

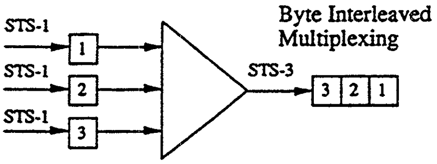

Byte Interleaved Multiplexing

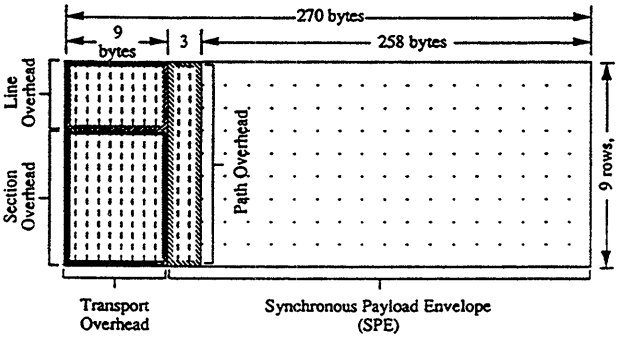

George H. Clapp

Standard CCITT SONET Rates
Synchronous Transport Module Levels

Supports "super-rate" services which require multiples of the STS-1 payload capacity, e.g., Broadband ISDN H_4 channel.

N STS-1s are concatenated into a single structure and transported as a single entity.

Standardized within CCITT as Synchronous Transport Modules Level N, or STM-N.

North American format referred to as STS-N "Concatenated" (STS-Nc); STS-3c ≈ STM-1.

STM-1 is the CCITT basic building block.

STM Rate	Line Rate (Mb/s)
STM-1	155.52
STM-4	622.08
STM-8†	1244.16
STM-12†	1866.24
STM-16†	2488.32

†Candidate rates which are not standardized.

George H. Clapp

CCITT STM-1 Format

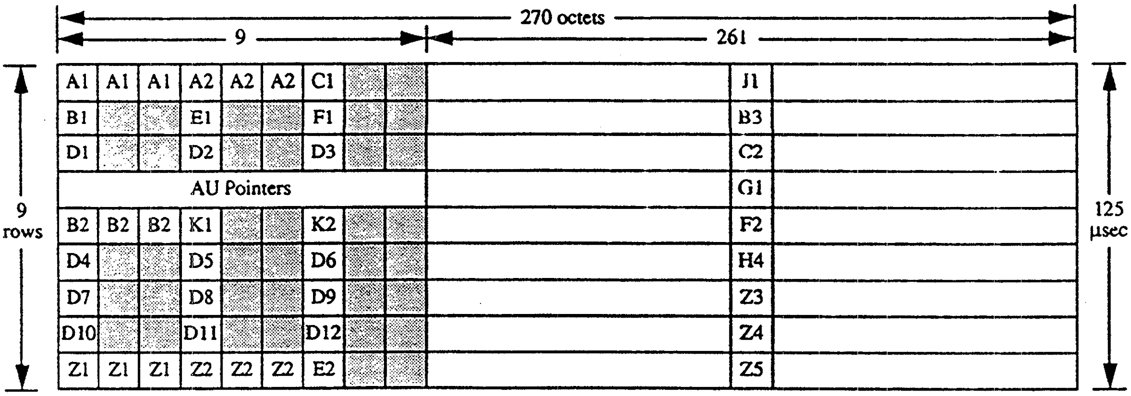

A1	A1	A1	A2	A2	A2	C1	▨	▨		J1	
B1	▨	▨	E1	▨	▨	F1	▨	▨		B3	
D1	▨	▨	D2	▨	▨	D3	▨	▨		C2	
AU Pointers										G1	
B2	B2	B2	K1	▨	▨	K2	▨	▨		F2	
D4	▨	▨	D5	▨	▨	D6	▨	▨		H4	
D7	▨	▨	D8	▨	▨	D9	▨	▨		Z3	
D10	▨	▨	D11	▨	▨	D12	▨	▨		Z4	
Z1	Z1	Z1	Z2	Z2	Z2	E2	▨	▨		Z5	

A1, A2	Framing	B3	BIP-8
AU Pointers	Administrative Unit Pointers	C2	STM Identifier
B1	BIP-8 (Bit Interleaved Parity)	F2	User-Defined Channel
B2	BIP-24	G1	Path Status
C1	STM Identifier	H4	Multiframe Indicator
D1-D12	Data Communications Channel	J1	Path Trace
E1, E2	Order Wire	Z3-Z5	Growth (reserved as spare)
F1	User-Defined Channels		
K1, K2	Automatic Protection Switching (APS)	▨	Reserved for national use
Z1-Z2	Growth (reserved as spare)		

George H. Clapp

SONET
Proposed STM-1/STS-3c Cell Mapping

Three mechanisms:

Pointer.

“Multiframe Indicator” octet (H4) of the Path Overhead points to the beginning of the first complete cell following the H4 octet.

Cell counting.

Header Check Sequence (HCS) calculation.

Following cell boundary identification via the H4 octet, the HCS is checked to verify cell delineation.

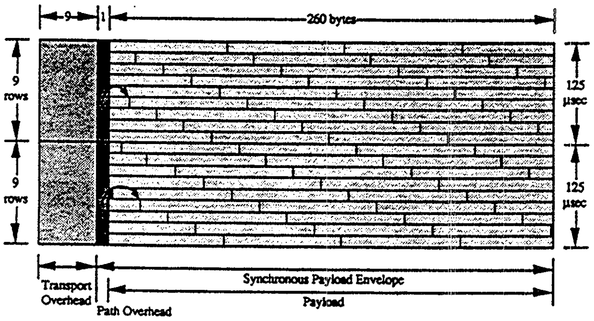

$$\frac{260 \text{ bytes} \times 8 \times 9}{125\ \mu\text{sec}} = 149.760 \text{ Mb/s} \qquad \frac{(260 \times 9) \text{ bytes / frame}}{53 \text{ bytes / cell}} = 44\frac{8}{53} \text{ cells / frame}$$

$$\frac{44\frac{8}{53} \text{ cells / frame} \times 48 \text{ bytes / cell} \times 8}{125\ \mu\text{sec}} \approx 135.632 \text{ Mb/s}$$

SONET
Proposed STM-1/STS-3c Cell Mapping

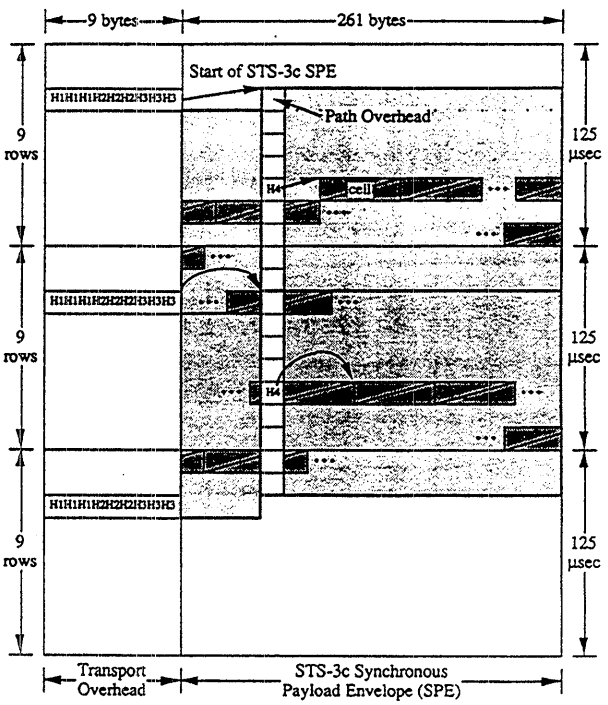

SONET
ATM Cell Delineation

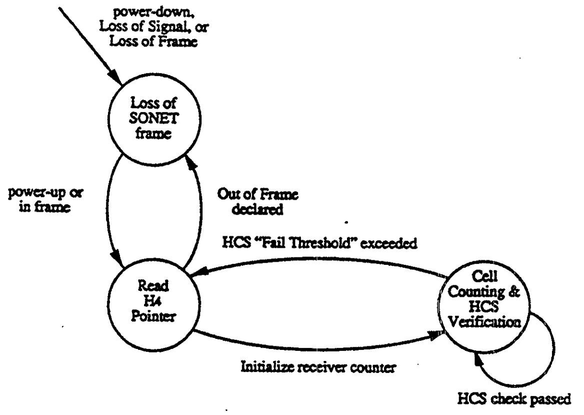
power-down, Loss of Signal, or Loss of Frame
Loss of SONET frame
power-up or in frame
Out of Frame declared
HCS “Fail Threshold” exceeded
Read H4 Pointer
Cell Counting & HCS Verification
Initialize receiver counter
HCS check passed

References

A SONET Cell Alignment Procedure Using the Header Check Sequence, R. C. Lau, ECSA Document T1S1.1/89-030, March 6, 1989.

An Analysis of Error Protection Mechanisms, M. Littlewood, IEEE Document P802.6-89/31, May 22-26, 1989, Perth, Western Australia, Australia.

ATM mapping for SDH, UNI structure, J. Anderson, ECSA Document T1S1.1/89-163, May 15-19, 1989, East Bruswick, NJ.

Digital Hierarchy Formats Specifications, American National Standard for Telecommunications, American National Standard T1.107-1988.

Distributed Queueing: Performance Characterisation, R. M. Newman, IEEE Document P802.6-88/11, March 13-17, 1988, San Diego, CA.

Engineering and Operations in the Bell System, R. F. Rey, Technical Editor, Bell Telephone Laboratories, 2nd Edition, 1983.

Error Performance of Data Adaptation Layer for B-ISDN, S. Dravida, ECSA Document T1S1.1/89-245, May 15-19, East Brunswick, NJ, USA.

Error Performance of Data Adaptation Layers for IEEE 802.6, S. Dravida, IEEE Document P802.6-88/100, November 4, 1988, Phoenix, AZ, USA.

Generic System Requirements in Support of Switched Multi-megabit Data Service, Bellcore Technical Advisory TA-TSY-000772, Issue 2, March 1989.

Improving DQDB Fairness, E. L. Hahne, N. F. Maxemchuk, A. K. Choudhury, IEEE Document 802.6-89/52.

Improving the Fairness of Distributed-Queue-Dual-Bus Networks, E. L. Hahne, A. K. Choudhury, N. F. Maxemchuk, Submitted to Infocom '90.

Local Access System Generic Requirements, Objectives, and Interface in Support of Switched Multi-megabit Data Service, Bellcore Technical Advisory TA-TSY-000773, Issue 1, December 1988.

Multiple Use of Cells in the IEEE 802.6 Dual Bus, A. Perdikaris, IEEE Document P802.6-88/10, March 13-18, 1988, San Diego, CA.

Proposed Standard Distributed Queue Dual Bus (DQDB) Metropolitan Area Network (MAN), P802.6/D6, March 3, 1989.

Scope and Functions of the ATM Adaptation Layer for Bursty Data Services, ECSA Document T1S1.1/89-243, May 15-19, 1989, East Brunswick, NJ.

T1S1 Techaical Sub-Committee Broadband Aspects of ISDN Baseline Document, R. Sinha, editor, ECSA Document T1S1.1/89-200, September 1989.

IEEE 802.6 Working Group chair
James F. Mollenauer
Prime Computer Inc.
500 Old Connecticut Path (MS 10C 27)
Framingham MA 01701
508-879-2960 x4085

For technical information concerning SMDS, contact:
Ram B. Misra
Bellcore
445 South Street Room 2K-244
Morristown, NJ 07960-1910
201-829-4128

T1S1.5 Broadband Subworking Group chair
Glenn Estes
Bell Communications Research
Morristown Research & Engineering
Room 2F-295
435 South Street
Morristown, NJ 07960
201-829-4238

For copies of IEEE 802 documents, contact:
Alpha Graphics
10215 North 35th Avenue
Phoenix, AZ 85051
602-863-0999

3.3.6 Point-to-Point Protocol Extentions (pppext)

Charter

Chairperson:

Stev Knowles, `stev@ftp.com`

Mailing Lists:

General Discussion: `ietf-ppp@ucdavis.edu`
To Subscribe: `ietf-ppp-request@ucdavis.edu`

Description of Working Group:

The Point-to-Point Protocol (PPP) was design to encapsulate multiple protocols. IP was the only network layer protocol defined in the original documents. The working group is defining the use of other network level protocols and options for PPP. The group will define the use of protocols including: bridging, ISO, DECNET (Phase IV and V), XNS, and others. The group will also define new PPP options for the existing protocol definitions, such as stronger authentication and encryption methods.

Goals and Milestones:

Aug 1990	The main objective of the working group is to produce an RFC or series of RFCs to define the use of other protocols on PPP.

CURRENT MEETING REPORT

Reported by Stev Knowles/ FTP

At the PSC meeting, the PPP working group decided to accept the current options draft, with a minor correction to the text. The correction removed the requirement to drop packet types that have not been defined. This was intended to allow people to use a PPP like protocol for things not covered by the document. Representatives from AT&T interested in Frame Relay technology asked for help in using PPP over Frame Relay circuits. The change they required should have been introduced into the process at a much earlier stage. Unfortunately, the AT&T people did not know about this process. The group decided to make this one change, that will allow the Frame Relay people to write an options document/ application note telling how to use PPP over Frame Relays.

Fred Baker (VitaLink) presented his final draft of the bridging document, and has made the final suggested changes.

Several people volunteered to help write specs, including:

Dave Katz	ISO
Heather Dean	Frame Relay
Art Harvey	ISO, DECNet IV
Steve Senum	Appletalk, DECNet IV
Larry Backman	IPX
Frank Kastenholz	MIB
Frank Slaughter	Appletalk
John Loverso	MIB
Stev Knowles	MIB

The next meeting of the group will be a video conference held June 18, 1990.

ATTENDEES

Fred Baker	baker@vitalink.com
Steve Bruniges	
Bruce Crabill	bruce@umdd.umd.edu
Stan Froyd	sfroyd@salt.acc.com
Russell Hobby	rdhobby@ucdavis.edu
Frank Kastenholz	kasten@interlan.interlan.com
David Kaufman	dek@proteon.com
Tony Lauck	lauck@dsmail.dec.com
John R Loverso	loverso@xylogics.com
Keith McCloghrie	sytek!kzm@hplabs.hp.com
Zbigniew Opalka	zopalka@bbn.com
Paul Parker	paul.parker@cs.cmu.edu
Mike Patton	map@lcs.mit.edu
Drew Perkins	ddp@andrew.cmu.edu
Steve Senum	sjs@network.com
Frank Slaughter	fgs@shiva.com
Richard Smith	smiddy@pluto.dss.com
Cal Thixton	cthixton@next.com

3.3.7 Router Discovery (rdisc)

Charter

Chairperson:

Steve Deering, deering@pescadero.stanford.edu

Mailing Lists:

General Discussion: gw-discovery@gregorio.stanford.edu
To Subscribe: gw-discovery-request@gregorio.stanford.edu

Description of Working Group:

The Router Discovery Working Group is chartered to adopt or develop a protocol that Internet hosts may use to dynamically discover the addresses of operational neighboring gateways. The group is expected to propose its chosen protocol as a standard for gateway discovery in the Internet.

The work of this group is distinguished from that of the Host Configuration Working Group in that this group is concerned with the dynamic tracking of router availability by hosts, rather that the initialization of various pieces of host state (which might include router addresses) at host-startup time.

Goals and Milestones:

Done Created working group; established and advertised mailing list. Initiated email discussion to identify existing and proposed protocols, for router discovery.

Done Held first meeting in Palo Alto. Reviewed 9 candidate protocols, and agreed on a hybrid of cisco's GDP and an ICMP extension proposed by Deering.

Done Held second meeting in Tallahassee. Reviewed the proposed protocol and discussed a number of open issues.

Done Held third meeting in Pittsburgh. Discussed and resolved several issues that had been raised by email since the last meeting. Draft specification of router discovery protocol to be ready by next meeting. Experimental implementations to be started.

Aug 1990 Meet in Vancouver. Review draft specification, and determine any needed revisions. Evaluate results of experimental implementations and assign responsibility for additional experiments, as required. Submit the specification for publication as a Proposed Standard shortly after the meeting.

Oct 1990 Revise specification as necessary, based on field experience. Ask the IESG to elevate the protocol to Draft Standard status. Disband.

CURRENT MEETING REPORT

Reported by Steve Deering/Stanford

Minutes:

This was the third meeting of the Router Discovery Working Group.

Steve Deering started by apologizing for not having produced a draft specification of the proposed router discovery protocol in time for this meeting, and promised to do so before the July IETF meeting.

He then reviewed the current proposal, for the sake of newcomers.

The router discovery protocol uses a pair of new ICMP messages:

- Router Advertisement messages, which are periodically multicast by routers.
- Router Solicitation messages, which may be multicast by hosts at start-up time only, to solicit immediate Router Advertisements instead of waiting for the periodic ones.

(These were formerly called "Router Report" and "Router Query" messages, respectively.)

Advertisements are sent to the "all-hosts" IP multicast address, and solicitations are sent to the "all-routers" IP multicast address. Hosts and/or routers may be configured to use IP broadcast addresses instead, though this is discouraged. The router discovery protocol is not applicable to networks that do not support either IP multicast or IP broadcast.

The format of the Router Advertisement message is as follows:

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |     Code      |           Checksum            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Reserved    |     Count     |         Holding Time          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Router Address                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Preference Level                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Router Address                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Preference Level                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               .                               |
|                               .                               |
|                               .                               |
```

Type — identifies this as an ICMP Router Advertisement.

Checksum — standard ICMP checksum.

Code and Reserved — zero.

Count — number of (Router Address, Preference Level) pairs included in the message.

Holding Time — number of seconds that hosts should consider the information in this message to be valid.

Router Address — one of the sending router's IP addresses on the physical network on which this message is sent.

Preference Level — preferability of the preceding router address as a default router address, relative to other router addresses belonging to the same IP subnet.

The usual case in which a router has more than one address on a single physical network is when that network is supporting more than one IP subnet. A receiving host is expected to ignore those (Router Address, Preference Level) pairs that do not belong to the same IP subnet as the host. (This implies that a host must know its own IP address and subnet mask before it may use the information in a Router Advertisement message.)

The format of the Router Solicitation message is as follows:

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |     Code      |           Checksum            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Reserved                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Type identifies this as an ICMP Router Solicitation.

Checksum standard ICMP checksum.

Code and Reserved zero.

The group then discussed a number of topics that had been raised on the mailing list since the previous meeting.

- Preference Level field

  Deering tried again to convince the group that the Preference Level field was unnecessary and undesirable, and again he failed. It was agreed that the field shall be present in the Router Advertisement messages, if for no other reason than that the Host Requirements document requires a preference level to be associated with each default router (even though it does not require hosts to do anything with it).

  Deering then proposed that the Preference Level be encoded as a signed, 32-bit, twos-complement integer, such that a higher value means more preferable. A router that is not configured with a specific preference level (or that does not compute its own preference level, based on routing information), will advertise a level of 0. The minimum level (80000000 hex) is reserved to indicate routers that must not be used as default routers (i.e., that may be used only for specific destinations, of which the hosts have been informed by ICMP Redirect or static configuration).

  Greg Satz had proposed that a router's preference level be derived from that router's metric for its "default" route, rather than from manual configuration. After some discussion of the merits and weaknesses of that approach, it was agreed that it would be allowed but not required by the router discovery specification. It was noted that a routing metric will normally have to be converted to a preference level, rather than being used directly, since for most routing metrics, smaller values mean more preferable.

No objections were raised to Deering's proposed encoding for the Preference Level field.

- multicast vs. unicast replies to Solicitations
- dallying

Two unresolved issues were: should Advertisements sent in response to Solicitations be multicast or unicast, and should randomized delays be required before Solicitations and/or before responding Advertisements? Some people felt that dallying before Solicitations was important to prevent massive collisions when a LANful of hosts all boot at once, for example, after power is switched on (in a classroom, say) or is restored after a power failure. After much debate, it was agreed that hosts should dally for a short, random interval (between 0 and 1 seconds was suggested) before sending a Solicitation. If a host receives an Advertisement while dallying, it should refrain from sending a Solicitation.

The optimal router behavior in response to a Solicitation is not at all clear – a case was made for dallying or not, and for either unicast or multicast responses. Therefore, this will be left to the implementors' discretion for now, with a suggestion that the behavior be configurable. The group would welcome any analysis, simulation results, or reports of field experience that might favor a particular behavior.

- periodic advertising rate

Another outstanding issue was how often the periodic, unsolicited Advertisements should be sent. The choice depends on whether or not the advertisements are being used for black-hole detection, in addition to simple router discovery. For black-hole detection, the advertising rate must be high enough to allow router failures to be detected before transport connections fail (an interval of 10 seconds is the value used for this purpose in the ISO ES-IS protocol). If the advertisements are only used for router discovery, a much longer interval (10 minutes, say) would be adequate – in this case the periodic advertisements serve only for recovery from the situation in which hosts and routers boot up on different sides of a subnet partition, which is later healed.

In the absence of agreement on how black-hole detection should be done, the advertising interval must be configurable. The initial version of the document will suggest a default interval of 10 minutes. Subsequent decisions on black-hole detection may cause a smaller default value to be recommended.

- black-hole detection

  Once the router discovery specification has been agreed upon, it has been suggested that this Working Group might turn its attention to the black-hole detection problem. A general discussion of the problems and possible solutions ensued, with no agreements being reached. (It was pretty much a rehash of previous discussions on the group mailing list; an archive of those messages is available for the interested reader.)

Action Items

- Deering to generate a draft Router Discovery specification before the July IETF meeting.
- Experimental implementations of the proposed protocol to be developed and deployed – no promises, but Andrew Cherenson and John Veizades both offered to help (presumably for Unix and for the Macintosh OS, respectively), as soon as Deering gets the specification done. Greg Satz has previously offered the source code for his BSD Unix implementation of cisco's Gateway Discovery Protocol (GDP) as one possible starting point.

Next Meeting

The Router Discovery Working Group will next meet in Vancouver, at the July/August IETF meeting.

ATTENDEES

Pat Barron	pat@transarc.com
Fred Bohle	fab@saturn.acc.com
Steven Bruniges	
David Burdelski	daveb@ftp.com
Duane Butler	dmb@network.com
John Cavanaugh	J.Cavanaugh@StPaul.NCR.COM
Andrew Cherenson	arc@sgi.com
Noel Chiappa	jnc@PTT.LCS.MIT.EDU
Steve Deering	deering@pescadero.stanford.edu
Dave Forster	
Richard Fox	sytek!rfox@sun.com
Karen Frisa	karen@kinetics.com
Steve Hubert	hubert@cac.washington.edu
Van Jacobson	van@helios.ee.lbl.gov
Stev Knowles	stev@ftp.com
Yoni Malachi	yoni@cs.stanford.edu
Keith McCloghrie	sytek!kzm@hplabs.HP.COM
Leo J. McLauglin III	ljm@twg.com
Jeff Mogul	mogul@decwrl.dec.com
John Moy	jmoy@proteon.com
Mike Patton	MAP@LCS.MIT.EDU
Drew Perkins	ddp@andrew.cmu.edu
Stephanie Price	cmcvax!price@hub.ucsb.edu
Michael Reilly	reilly@nsl.dec.com
Greg Staz	satz@cisco.com
Tim Seaver	tas@mcnc.org
Frank Slaughter	fgs@shiva.com
Richard Smith	smiddy@pluto.dss.com
Brad Strand	bstrand@cray.com
Cal Thixton	cthixton@next.com
John Veizades	veizades@apple.com
Jonathan Wenocur	jhw@shiva.com

3.3.8 Router Requirements (rreq)

Charter

Chairperson:

Jim Forster, forster@cisco.com
Philip Almquist, almquist@jessica.stanford.edu

Mailing Lists:

General Discussion: ietf-rreq@Jessica.Stanford.edu
To Subscribe: ietf-rreq-request@Jessica.Stanford.edu

Description of Working Group:

The Router Requirements Working Group has the goal of rewriting the existing Router Requirements RFC, RFC-1009, and a) bringing it up to the organizational and requirement explicitness levels of the Host Requirements RFC's, as well as b) including references to more recent work, such as the RIP RFC and others.

Goals and Milestones:

May 1990	Produce a draft document for initial comment by the community.

CURRENT MEETING REPORT

Reported by Vince Fuller/Stanford and Philip Almquist/ Consultant

Minutes: Tuesday, 1-May (AM)

- TOS routing
  - Two aspects - internal queue handling vs. next hop choice
    - RREQ document deals primarily with later (external behavior)
  - Number of bits: 3 currently defined for TOS, 2 other "spare" bits
    - does router need to know what bits mean or can it just match against information available via routing protocols?
    - is TOS a hint or a requirement (more discussion later) - hint implies it is safe to ignore extra bits for now
    - issue: other groups may want to use those two bits for other things
    - requirement: all routers must make the same routing choices regarding TOS and all must implement TOS (but not all protocols will use) to prevent routing loops (Chair's statement)
    - issue: may have to change routing protocols if the number of bits changes (tough)
    - quote: "keep your paws off those two bits" - JNC, Area Director
    - DECISION: use 3 bits (problem made moot by above quote)
- TOS semantics:
  - hint philosophy: deliver packets to default TOS if no match exists
  - requirement philosophy: drop packets if no TOS match exists (editors note: a very long and heated discussion of these differing philosophies consumed most of the first part of this session)
  - TOS unreachable ICMP message for "requirement" case
  - Proteon OSPF implementation allows per-TOS metric setting on lines; setting to infinity and dropping on no TOS match allows some small amount of policy control over line usage
  - problem: handling of TOS unreachable message is undefined
  - problem: won't work if host ignores TOS unreachables (or falls-back automatically to default TOS)
  - problem: TOS bits are not defined absolutely (i.e. in bps, etc.)
  - suggestion (Satz): two sides write up their cases; include both in draft document for further review (any takers?)
  - idea: use one of the "unused" bits to specify hint/required TOS
  - Milo believes looping can occur if TOS is treated as a hint - need specific scenarios

- Fragmentation
  - Review of each option outlined in draft text:
    - option 1: no longer needed by MTU discovery
    - option 3: invalid, since 576 is NOT the minimum required MTU size
  - if anyone has empirical evidence of a good way to do it, them document should discuss it; otherwise, defer to IP RFC
    - action: talk to Jeff Mogul and Van Jacobson about their experiments
- Reassembly
  - Router MUST reassemble packets destined to itself (i.e., ICMP messages)
    - router is acting as host in this case, must follow HR
    - HR says must reassemble max of 576 bytes or connected interface MTUs
  - Router MUST NOT reassemble packets that are forwarded
    - reassembly not possible if multiple paths exist, etc.
  - Multicast handling
    Minimal discussion. Steve Deering ("Mr. Multicast") volunteered to write some draft text
- TTL
  - Long discussion about schizophrenic use of TTL as time AND hop count
    - TCP makes assumptions about real time of packet life vs. TTL handling (problem can occur with sequent number wraparound)
    - no implementors expect to implement use of TTL as time (fact of life)
    - deprecate "SHOULD" to "MAY" for decrementing TTL by time
    - include discussion of why this should be done (what TCP expects, etc.)
  - Handling of TTL boundary conditions:
    - TTL 0 - router MUST NOT drop packets to itself on TTL = 0 (HR sez)
    - router MUST NOT ever send a packet with TTL = 0 (ditto)
    - router SHOULD return ICMP time exceeded if it decrements TTL to 0
    - router MUST NOT "pre-discard" packets with TTL > 0 even if it knows (via link-state routing, for example) how many hops a destination is (it breaks "traceroute" to do so and doesn't really gain much)
    - should there be some discussion of "traceroute's" expectations?

Minutes: Tuesday, 1-May (PM)

A review of writing assignments/document sections was done. The following (mostly un-written) sections were worthy of mention (mainly because no one is doing anything about them):

7. Routing protocols - RIP, EGP, BGP (Yakov Rekhter/IBM), OSPF

8. Network Management Protocols - SNMP (Steve Willis/Wellfleet), CMIP/CMOT

9. Administrative and Policy controls, including: filters (both traffic and routing info), interchange between EGP's and IGP's, preference of routes by [protocol, neighbor, network number, etc.], conditions for default generation, etc. (subcommittee formed and had preliminary discussion over lunch - included Steve Willis/Wellfleet, Philip Almquist/Consultant/Stanford, Vince Fuller/Stanford/BARRNet, Michael Reilly/DEC, and one or two others forgotten by the editor – they and others interested in these issues (Milo, Jeff?) should get in touch with the Chair ASAP)

10. Initialization, operation, management

Appendix A - Internet-specific requirements

Appendix B - Requirements for specific uses (i.e., regional network)

- Multicast
  - forwarding of multicasts is not yet required
  - router SHOULD perform host multicast functions (per RFC1112 and HR)
  - router MUST NOT pass "letter-bomb" multicasts (as target of source route)
  - record route doesn't present a problem (according to Steve D.)
  - multicasts should not be used as a hop in a source route either
- TOS, take 2 (Yakov had a few things to say)
  - in current Internet, virtually no use (according to NSFNet statistics)
  - chicken and egg problem (Steve D.)
  - 3 bits are too coarse to be useful for policy control
  - all routers in AS (at least) must make same routing decision on TOS in order to prevent loops
    - what about for paths through multiple AS's?
    - what if AS's are multi-homed?
  - how to use in presence of sources (protocols) of routing information?
    - use to prefer protocol if it has exact TOS match?
  - opinion: TOS 0 is default - must always exist and is used if no exact match
  - opinion: forbid setting of multiple TOS bits ("Christmas tree") - enforce by treating as TOS 0 (?)

  - no definite conclusion (no surprise here!)

- Broadcast handling
  - Directed broadcasts
    - routers MUST support, MAY provide knob to disable
    - justification: widely used, part of IP architecture
  - All-subnets broadcast
    - current behavior: only sent to first subnet seen
    - Chair will make case to IESG to make it an obsolete part of the IP architecture (by creating a successor to RFC922)
    - consider as a SHOULD NOT - may support, but MUST provide knob which defaults behavior to disabled
- IP options
  - Record route - MAY in HR, specify as MUST in RREQ
  - Timestamp - ditto
    - long discussion about when during packet processing the timestamp should be added - no conclusion
    - document should state that when it happens is not defined and will be implementation-dependent
    - Yakov (opinion): all routers should do timestamp at same point in packet processing - not much agreement from rest of WG
  - Option insertion by routers
    - security option must be inserted, so it MUST be allowed (RFC1108)
    - what if no option space available - Martin Gross/DCA will address
    - are there other options that need to be inserted?
  - non-understood options - MUST be passed unchanged
  - source route - only one source route option may exist
- Precedence
  - OSPF mandates that routers set precedence to Internet Control
  - BGP - ditto
  - issue: may be political problems with this
  - what about network management traffic?
  - DCA group is working on paper describing scheme
- Martian address filtering
  MUST provide functionality and provide switch to enable/disable (long discussion ensued about performance impact of making it strictly a MUST)
- What's next?
  - video-conference will take place in June (tentatively, Monday, June 11th)
  - Internet-Draft is expected after August IETF

ATTENDEES

Douglas Bagnall	bagnall_d@apollo.hp.com
Pablo Brenner	
Steven Bruniars	
John Cavanaugh	john.cavanaugh@stpaul.ncr.com
Curtis Cox	zk0001@nhis.navy.mil
Steve Deering	deering@pescadero.stanford.edu
Dino Farinacci	dino@bridge2.3com.com
Dave Forster	
Karen Frisa	karen@kinetics.com
Stan Froyd	sfroyd@salt.acc.com
Vince Fuller	fuller@jessica.stanford.edu
Martin Gross	martin@protolaba.dca.mil
Keith Hogan	keith%penril@uunet.uu.net
Kathy Huber	khuber@bbn.com
Steve Hubert	hubert@cac.washington.edu
Tim Hunter	thunter@allegum.bitnet
Phil Karn	Karn@Thumper.Bellcore.Com
Frank Kastenholz	kasten@interlan.interlan.com
Stev Knowles	stev@ftp.com
Kanchei Loa	loa@sps.mot.com
Yoni Malachi	yoni@cs.stanford.edu
Milo Medin	medin@nsipo.nasa.gov
David Miller	dtm@mitre.org
John Moy	jmoy@proteon.com
Phil Park	ppark@bbn.com
Drew Perkins	ddp@andrew.cmu.edu
Paul Pomes	paul-pomes@uiuc.edu
Stephanie Price	price@cmcvax!uscbcsl.edu
Stepanie Price	price@mcvax!ucsbcsl.edu
Michael Reilly	reilly@nsl.dec.com
Joel Replogle	replogle@ncsa.uiuc.edu
Greg Satz	satz@cisco.com
Steven Senum	sjs@network.com
Jim Sheridan	jsherida@ibm.com
Richard Smith	smiddy@dss.com
Steve Storch	sstorch@bbn.com
Roxanne Streeter	streeter@nsipo.arc.nasa.go
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Kannan Varadhan	kannan@oar.net
John Veizades	veizades@apple.com

Peter Vinsel	farcomp!pcv@apple.com
John Vollbrecht	jrv@merit.edu
David Waitzman	djw@bbn.com
Y Wang	
Jonathan Wenocur	jhw@shiva.com
Steve Willis	swillis@wellfleet.com
Walt Wimer	ww0n@andrew.cmu.edu
John Wobus	jmwobus@suvm.acs.syr.edu
Richard Woundy	rwoundy@ibm.com
Sze-Ying Wuu	wuu@nisc.junc.net
Mary Youssef	mary@ibm.com

3.4 Network Management Area

Director: Dave Crocker/DEC

Area Summary

Reported by Greg Vaudreuil /CNRI

The Network Management area currently has 10 active working groups. Of those groups the Alert Management, Decnet Phase IV MIB, the FDDI MIB and Transmission MIB, Call Accounting, Management Services Interface and the OSI Internet Management Working Groups met.

The Alert Management Working Group has completed their specifications and they will be submitted to the IESG for consideration at the August Plenary meeting. The OIM Working Group is expected to present their latest CMIP-over-TCP document to the IETF and IESG at the August IETF meeting.

3.4.1 Alert Management (alertman)

Charter

Chairperson:

Louis Steinberg, louiss@ibm.com

Mailing Lists:

General Discussion: alert-man@merit.edu
To Subscribe: alert-man-request@merit.edu

Description of Working Group:

The Alert Management Working Group is chartered with defining and developing techniques to manage the flow of asynchronously generated information between a manager (NOC) and its remote managed entities. The output of this group should be fully compatible with the letter and spirit of SNMP (RFC 1067) and CMOT (RFC 1095).

Goals and Milestones:

Done	Develop, implement, and test protocols and mechanisms to prevent a managed entity from burdening a manager with an unreasonable amount of unexpected network management information. This will focus on controlling mechanisms once the information has been generated by a remote device.
Done	Write an RFC detailing the above, including examples of its conforment use with both SNMP traps and CMOT events.
May 1990	Develop, implement, and test mechanisms to prevent a managed entity from generating locally an excess of alerts to be controlled. This system will focus on how a protocol or MIB object might internally prevent itself from generating an unreasonable amount of information.
Dec 1990	Write an RFC detailing the above. Since the implementation of these mechanisms is protocol dependent, the goal of this RFC would be to offer guidance only. It would request a status of "optional".

CURRENT MEETING REPORT

Reported by Lou Steinberg/ IBM

The Alert Management Working Group met at the Pittsburgh IETF in the context of a non-technical "Birds of a Feather" meeting. The meeting length was reduced in order to allow interested WG members the opportunity to attend an OIM Working Group session scheduled for the same time.

While the work of this group on the first "flow control" document has been officially completed, Dave Crocker asked that the shortened session be used to afford a final opportunity for interested parties to comment on the DRAFT. It was his opinion that several individuals present had not previously been able to give public comment. As there were no major concerns, and no implementation experience that contradicts the findings of current implementations, the remainder of the time was spent describing the document details to several new members who had not yet read the DRAFT.

The meeting was adjourned after John Cook repeated earlier calls for information from *any* vendor implementing alerts. John is currently authoring the second "informational" DRAFT on techniques for generating alerts.

ATTENDEES

Hussein Alaee	hussein_alaee@3mail.3com.com
Douglas Bagnall	bagnall_d@apollo.hp.com
Pablo Brenner	sparta!pbrenner@uunet
Ted Brunner	tob@thumper.bellore.com
Y C Wang	21040 Homestead Rd Cupertino, CA 24306
Jeffrey Case	case@utkux1.utk.edu
Martina Chan	mchan@cs.utk.edu
John Cook	cook@chipcom.com
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Frank Feather	feather@ece.cmu.edu
Metin Feridun	mferidun@bbn.com
Stanley Froyd	sfroyd@salt.acc.com
Ella Gardner	epg@gateway.mitre.org
Brian Handspicker	bd@vines.dec.com
Richard Hart	hart@decvax.dec.com
Frank Kastenholt	kasten@interlan.interlan.com
Tony Lauck	lauck@dsmail.dec.com
Roy Maxion	maxion@cs.cmu.edu
Keith McCloghrie	sytek!kzm@hplabs.hp.com
Greg Minshall	minshall@kinetics.kinetics.com

Oscar Newkerk	newkerk@decwet.dec.com
John O'Hara	saperia@tcpjon.enet.dec.com
Dave Perkins	dave_perkins@3mail.3com.com
Ron Roberts	roberts@jessica.stanford.edu
Jim Sheridan	jsherida@ibm.com
Rob Shirey	shirey@mitre.org
Mark Sleeper	
Richard Smith	smiddy@dss.com
Sudhanshu Verma	verma@hpindbu.hp.com
David Waituman	djw@bbn.com
Richard Woundy	rwoundy@ibm.com

3.4.2 Internet Accounting (acct)

Charter

Chairperson:
Cyndi Mills, cmills@bbn.com

Mailing Lists:
General Discussion: accounting-wg@bbn.com
To Subscribe: accounting-wg-request@bbn.com

Description of Working Group:

The Internet Accounting Working Group has the goal of producing standards for the generation of accounting data within the Internet that can be used to support a wide range of management and cost allocation policies. The introduction of a common set of tools and interpretations should ease the implementation of organizational policies for Internet components and make them more equitable in a multi-vendor environment.

In the following accounting model, this Working Group is primarily concerned with defining standards for the Meter function and recommending protocols for the Collector function. Individual accounting applications (billing applications) and organizational policies will not be addressed, although examples should be provided.

Meter <-> Collector <-> Application <-> Policy

First, examine a wide range of existing and hypothetical policies to understand what set of information is required to satisfy usage reporting requirements. Next, evaluate existing mechanisms to generate this information and define the specifications of each accounting parameter to be generated. Determine the requirements for local storage and how parameters may be aggregated. Recommend a data collection protocol and internal formats for processing by accounting applications.

This will result in an Internet draft suitable for experimental verification and implementation.

In parallel with the definition of the draft standard, develop a suite of test scenarios to verify the model. Identify candidates for prototyping and implementation.

Goals and Milestones:

May 1990	Policy Models Examined
Aug 1990	Meter Working Draft Written
Nov 1990	Collection Protocols Working Papers Written
Feb 1991	Meter Final Draft Submitted
Feb 1991	Collection Protocol Working Papers Reviewed
May 1991	Collection Protocol Recommendation

CURRENT MEETING REPORT

Reported by Cyndi Mills/ BBN Notes taken by Don Hirsh/ Meridian TC

Summary

This was the first meeting of the Internet Accounting Working Group. We outlined a hierarchical architecture for accounting within routers and discussed the types of meters to be used at each level.

Agenda

- Accounting Architecture
- Technical Reports
  - Internet Accounting Model
  - Liaison Activities (ANTF, OSI)
- Open Discussion
- Working Group Administration
  - Review Charter & Minutes
  - Identify and Assign Action Items

ACCOUNTING ARCHITECTURE

Due to performance constraints and the explosion in complexity, we believe that it is not practical to perform detailed accounting to the user-id level within all networks. [Ed. The reasons should be documented in the Meter Services Document.]

Therefore we identified 4 levels of accounting interest/architecture:

```
Backbones/National ---------------------------
                          /        \
Regional          ------------ --------------
                      /  \  \ /  /   \
Stub/Enterprise     ---  ---  ---  ---  ---

Host
```

Note that mesh architectures can also be built out of these components. Each network performs accounting functions for its immediate subscribers / connections. Subscribers come in two flavors - subscriber networks and subscriber hosts (end-users from the networking perspective).

We define backbone networks as bulk carriers that have only other networks as subscribers. Individual hosts are not directly connected to a backbone.

Backbones and regionals are closely related, and differ only in size, the number of networks connected via each port, and "geographical" coverage. Smaller Regionals may also have a few directly connected hosts, acting as hybrid regional/stub networks. We consider a regional network as a subscriber network to the backbone.

Stub networks have hosts as direct connects, although they may be combined by Enterprise networks treated in the same fashion as stub networks. For the stub/enterprise network provider, hosts are the end-users, the accountable entities. For the stub/enterprise network provider, host addresses are the finest-granularity accountable entities available at the IP level.

Hosts are ultimately responsible for identifying the end user. This information may be shared with the network, but it is the host's responsibility to do so. Host accounting is not discussed in detail, since homogeneous Internet Accounting is most practical at the network provider level, and should be performed within the network routers under the control of the network provider. (After all, the host is the customer, and if I were selling network services I'm not sure I'd rely on the customer to tell me how much he owes without having a mechanism to keep the customer honest...) In addition, implementing accounting in the routers spares us from requiring that each host variation (various hardware platforms and operating system versions) retrofit TCP/IP implementations to include accounting as a condition for being attached to a network which relies on accounting information.

ENTITIES: Each of the higher-level network (backbone and regional) account for two sets of entities - one set corresponds to the network's immediate subscribers and a parallel set (optional?) covers the subscribers of the network below. This two-tiered system enables:

- verification between provider and subscriber
- reconstruction of accounting information around a single transit network which does not perform accounting functions.

Backbone Level Entities:	Adjacent Network (Port),	Source/Dest Net Number
Regional Level Entities:	Src/Dest Network Number,	Src/Dest Host Adr
Stub Level Entities:	Source/Dest Host Address (End-user ID pair optional)	
Host Level Entities:	Operating System dependent.	Use OS accounting.

This allocation of accountable entities to network levels bears further examination. Particularly, it is important to understand what complexity accounting introduces at each network level.

Backbone Level Complexity: Collects by port ID, and can further subdivide by network numbers from the IP address.

Regional Level Complexity: Collects by host address pair only, since network numbers can be derived from the host traffic matrix off-line.

Stub Complexity: Collect host address pair in any case. Approaches:

1. Leave all else up to the local stub network and proprietary means if further information is required.
2. Define IP option containing accounting information.
3. Piggyback on the policy-based routing option and recommend how to use it.

Note on including destination addresses in the entity identifier: Maintaining a traffic matrix at all levels seems to be a fair amount of overhead, but destination information is required so often that it seems reasonable to include it. This way policy arrangements about who is billed for communicating pairs can be independent of the originator of the traffic.

SUB-ENTITIES: If we are aggregating information, the counts attributed to a single entity may need sub-categories. Suggested sub-categories are:

- protocol type
- quality of service
- types of counts

TYPES OF COUNTS: All networks count both packets and bytes for the accountable entities.

TIME-OF-DAY: We need to be able to register start and stop times of flows. These trigger times should automatically start new aggregations for the affected aggregate meters (i.e., cause meters to send their data along with the start and end times, and restart the meter at 0.).

QUALITY OF SERVICE: Unresolved. What quality of service items should we be able to specify? QOS distinctions come in many forms. For services such as throughput, reliability, and delay there is a question of how detailed the information should be regarding:

- What level of service was requested
- What level of service was offered (negotiated)
- What level of service was actually provided (considering outages, etc.)

Technical Reports

1. INTERNET ACCOUNTING MODEL
   See attached slides
2. LIASON ACTIVITIES
   The ANSI Accounting WG has OSI Accounting Drafts available.
   Report on April Autonomous Network Task Force (ANTF) Meeting on Internet Billing:
   - Billing Models discussed:
     - Fixed Fee
     - Usage Sensitive Billing
     - Quality of Service Sensitive Billing
     - Quotas
     - Subsidy Issues
     - Campus/Stub AD aggregate vs. end-user feedback
   - Issues raised:
     - High speed counting
     - Fraud
     - Credit limits
     - Cooperation between stub and backbone networks
     - How heterogeneous can the models be?
     - Interaction with congestion control, access control, routing
   - Liaison Activities
     - IETF Internet Accounting
     - SMDS Accounting
     - OSI Accounting
   - Suggested Experiments
     - Flow-based instrumentation (use this to identify and play with flows)
     - Resource reservation (We should suggest ST-2 or MacHip, a St. Louis sponsored entry)
     - Instrument applications to provide feedback window (have a window with a * amount to meter applications)
3. OPEN DISCUSSION

END-USER FEEDBACK - Can end-users influence policy? How about the ability to provide accounting feedback mechanisms to network users real-time as they use it, what that might look like and so forth? [Ed. This is somewhat orthogonal to the group charter at the application level, but was an interesting discussion worth keeping in mind.]

POLICY-BASED ROUTING - Their relation to us is in their use of the IP header's options field. We might put in a Kerberos-style identifier that associates a particular machine/user/virtual circuit with a unique token. This scheme might work between adjacent networks to track FLOWS though them, but would be difficult (!!!) to pull off on an internet-wide basis. Some one or two of us should attend the policy-based routing sessions regularly since they're working on similar problems. Negotiating quality of service is in the province of policy-based routing?

GRANULARITY OF DISTINGUISHABLE ENTITY - Two positions were discussed:

(a) IP-based accounting with only existing IP header information is sufficient.
(b) One should try to accommodate users and perform accounting by the user-id where it is feasible.

IDEAS ON IDENTIFYING THE END-USER TO THE ACCOUNTING MECHANISM

(a) PARSING TCP and APPLICATION LAYER PROTOCOLS FOR USER INFORMATION? What about parsing more than the IP header? Considered untenable in a router. Even if we dismiss the violation of protocol layering as "purist", we still must contend with higher processing overhead. Hosts would still need to be modified to ensure that the user information is present. But passive "watchers" like scopes could be employed on LANs.

(b) MODIFY THE IP HEADER TO ADD ACCOUNTING INFORMATION? We don't believe it will get implemented by all existing hosts. (i.e., practically impossible).

(c) USE IP OPTIONS? Router perspective: putting user-based accounting stuff in a router is too much processing overhead. Counter-Example: Tymnet billing is on a per-user id. Compromise: At a minimum, an IP packet that has user-level accounting information might be afforded a lower priority in the router's processing queue.

(d) VANISHING OPTIONS? Vollbrecht points out that router-to-router accounting and ES - IS accounting are separable problems. This led to a discussion of how to leave the user-id option in for the stub network's use and strip it from the header when sending the packet to the regional to reduce regional overhead. Still a performance issue, and what about the checksum? This should be investigated more thoroughly.

SERVERS - How does one account for mail that explodes at a list server? Is it the responsibility of the host, the list, or the person who sends to the server?

OSI ACCOUNTING - Since they are not defining meters yet, we will probably influence the OSI standard with our choices.

ADMINISTRATIVE DETAILS

Review of Charter

- Examine existing and hypothetical policies to understand what set of information is required to satisfy usage reporting requirements.
- Define specifications of accounting meters, local storage requirements, and aggregation granularities.
- Recommend a data collection protocol and representations for processingby accounting applications.
- Develop test scenarios to verify model.
- Guess we have to recommend mechanisms for formulating policy, though we don't want to sink in the policy swamp. Also need to consider implementation issues since they are practical and affect the "reasonableness" of recommendations.

Internet Accounting Action Items

Can we live with the proposed schedule? Sure.

The following areas should be addressed in preparation for the August 1990 IETF meeting except where otherwise noted.

- Outline of Meter Service Document => C.Mills
- Architecture Discussion => Mailing List
  - Levels of Metering (Do we have the right model?)
  - Define Meters
    - Entities (Done. Review only.)
    - Quantities (Done. Review only.)
    - Time of Day (Further development.)
    - Quality of Service (How to approach this?)
- Liaison Activities
  - ANTF => Z.Su
  - OSI Accounting => B.Handspicker, M.Seger
  - SMDS => Z.Su
- Explanation of Concepts (writeups due to mailing list)
  - R.Reschly suggested that accounting on a backbone is the integral of bandwidth utilization and that proportional utilization rather than

  absolute measure is a useful measure.
  - J.Galvin proposed to write up some of the discussion.
  - M.Roselinsky will expand upon the user-id/cookie for the IP options field.
  - C.Mills will summarize the applicability of policy-based to accounting.
  - D.Hirsh will summarize current policy/practice in the Internet community (e.g., digest the FARnet study, summarize BBN/SRI activity, etc.) in light of the proposal for meters. (First step towards test scenarios.)
- Unassigned Tasks (may be deferred) => Mailing List
  - Define Accounting Log Formats
    - Local Storage Requirements
    - Compatibility with Existing Protocols
  - Develop Testbed/Prototypes

ATTENDEES

Peblo Brenner	sparte!pbrenner@uunet
Martina Chan	mchan@mot.com
James Galvin	galvin@tis.com
Don Hirsh	hirsh@magic.meridiantc.com
Keith McCloghrie	sytek!kzm@hplabs.hp.com
Robert Reschly	reschly@brl.mil
Milt Roselinsky	cmcvax!milt@hvb.vcsb.edu
Mark Seger	seger@mjs1.ogo.dec.com
Brad Strand	bstrand@cray.com
Zaw-Sing Su	zsu@tsca.istc.sri.com
John Vollbrecht	jrv@merit.edu

Purpose of Accounting

Accounting -

1. the art or system of keeping and analyzing financial records
2. an explanation oſ one's behavior

(merriam-webster)

Feedback

- **Understand / influence user behavior**
- **Measure compliance with network policies**

Financial

- **Allocate costs based on usage**
- **Generate revenue**

Recommend starting with feedback only to provide baseline for understanding effects of tariffs on user behavior.

An Internet Accounting Model

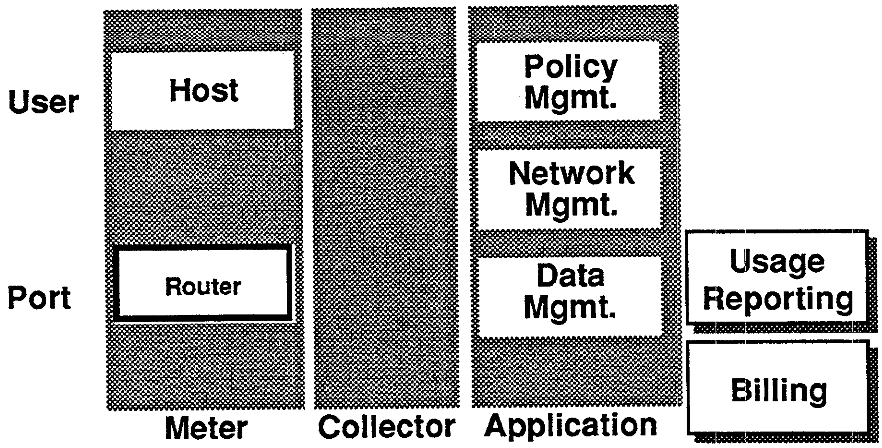

(analogous to OSI Accounting Model)

Locate Meters in the Router

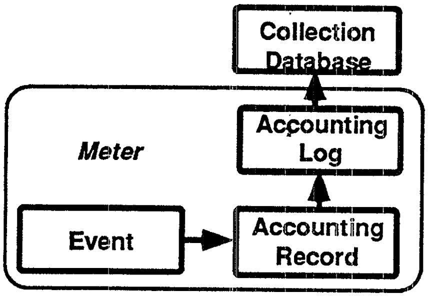

Motivation

- Minimize number of meters to reduce collection overhead.
- Meters located within "own" network assets are easiest to modify and control.

Implications of meters in the router

- Host systems must perform their own accounting *or*
- Hosts supply user/project ids to the network *and* routers use these ids for accounting.

Minimize Overhead by Metering at the IP Level

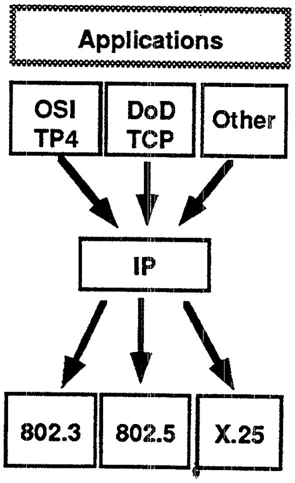

Protocol "Hourglass"

- Readily Accessible Information
  - –Hosts: Application, Transport, IP
  - IP Switches: Internet Only
- Minimize implementation and maintenance cost by minimizing the number of
  - –protocols
  - –vendors
  - –administrations
  - –physical units

affected.

Use Information Available in the IP Header

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version|  IHL  |Type of Service|          Total Length         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Identification        |Flags|      Fragment Offset    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Time to Live |    Protocol   |         Header Checksum       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Source Address                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Destination Address                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Options                    |    Padding    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Options may include: Policy-based flow, Security labels, special user ids, to-be-defined

Distinguish Entities Based on the IP Header

- **Router Port**
  - – identifies adjacent network only
- **IP Header**
  - **–Community**
    - the two unused bits in TOS
  - **–IP Address**
    - source/destination network number
    - source/destination IP address pair
  - **–IP Options**
    - User/Project IDs - requires host cooperation
    - Flow Identifier (like policy-based routing)
    - Policy Tags (e.g. reverse charging, non-interference)

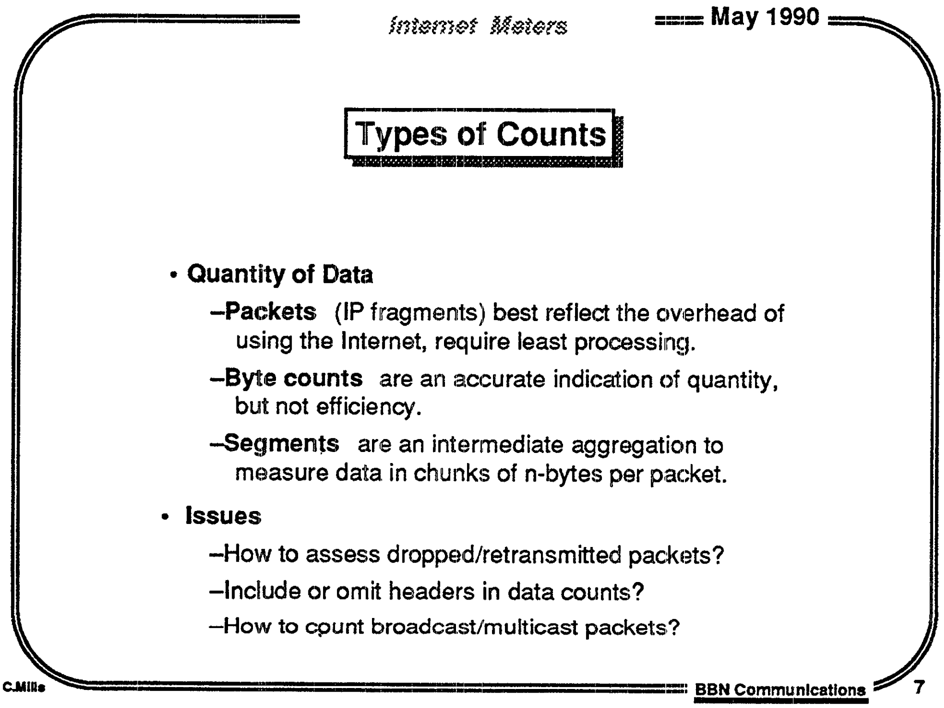

Types of Counts

- **Quantity of Data**
  - –**Packets** (IP fragments) best reflect the overhead of using the Internet, require least processing.
  - –**Byte counts** are an accurate indication of quantity, but not efficiency.
  - –**Segments** are an intermediate aggregation to measure data in chunks of n-bytes per packet.
- **Issues**
  - –How to assess dropped/retransmitted packets?
  - –Include or omit headers in data counts?
  - –How to count broadcast/multicast packets?

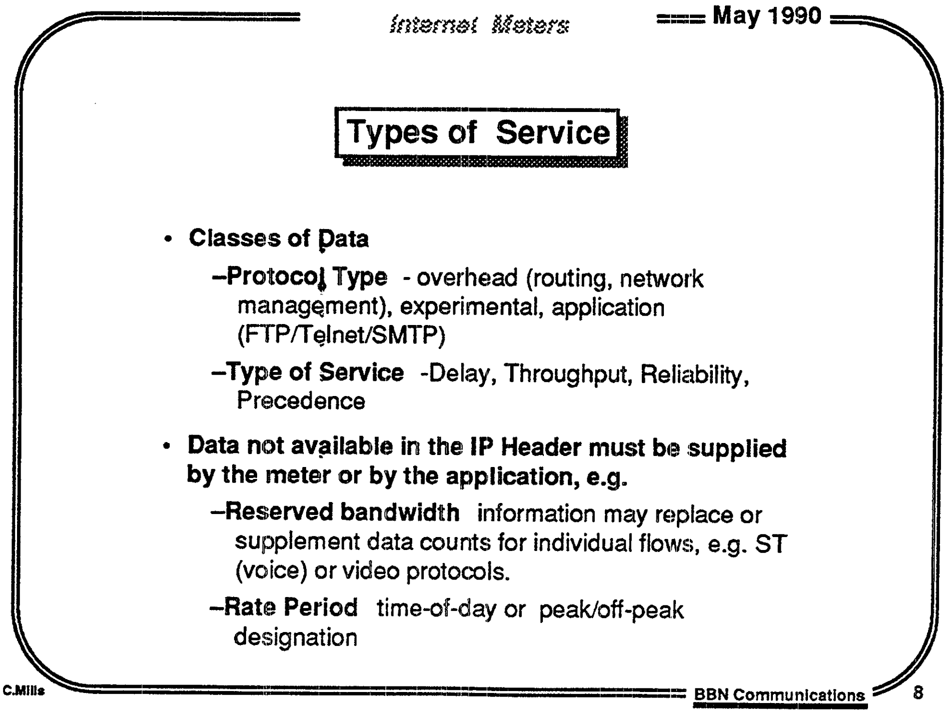

Types of Service

- **Classes of Data**
  - –**Protocol Type** - overhead (routing, network management), experimental, application (FTP/Telnet/SMTP)
  - –**Type of Service** -Delay, Throughput, Reliability, Precedence
- **Data not available in the IP Header must be supplied by the meter or by the application, e.g.**
  - –**Reserved bandwidth** information may replace or supplement data counts for individual flows, e.g. ST (voice) or video protocols.
  - –**Rate Period** time-of-day or peak/off-peak designation

Accounting Record Definition

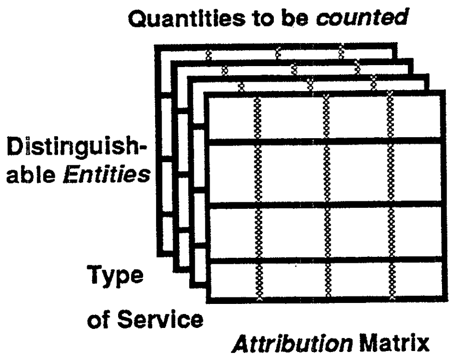

- Three dimensions
  - –Entities (user and/or subscriber)
  - –Types of Counts (packets, bytes)
  - –Types of Service (protocol id, TOS bits)
- Two dimensions is easier
  - –Entity (subscriber, user, type of service
  - –Counts

Attribution Granularity

Support complex decisions in the meter or
Apply policy retroactively at the application.

- Simple Attribution
  - –Attribute counts to single distinguishable entities.
- Community of Interest
  - –use static tables and algorithms (like Mailbridges) to assign combination of IP attributes to a single community of interest, reducing number of "buckets".
- Host Traffic Matrix
  - –specific subset of flow-based algorithm, based on combination of IP addresses
- Flow Based Decisions
  - –use dynamic servers and/or algorithms
  - –most flexible - can establish flow identifier on any combination of IP attributes and server data
  - –policy-based routing flow may not provide sufficient granularity

Servers

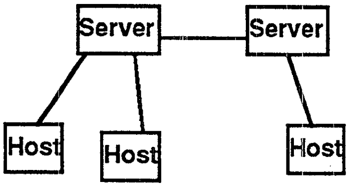

- Examples:
  - –Electronic Mail Applications and Relays
  - –OSI Application Gateways

- Services performed on behalf of the user may be performed at remote points in the network and performed for aggregation of users.
  - –Servers perform separate accounting for services which passes network expenses back to the user

  or

  - –Servers use IP tags to identify actual users

Backbone vs. Stub Accounting

Problem

- Accounting for every user in every network at every router is too fine a granularity - too much overhead.

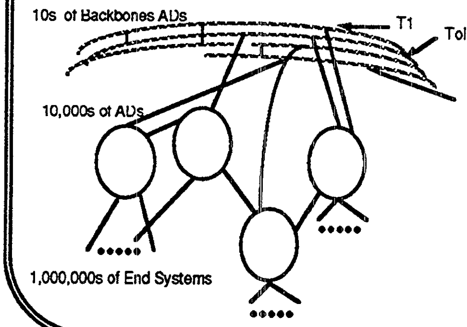

Approaches to Reducing Overhead

- Use coarser granularity for measuring transit traffic and/or in backbones.
- Hold each network responsible / accountable for its own traffic.
- Perform accounting only at entry and/or exit gateways.
- Use complex granularities for exception traffic only.

3.4.3 DECnet Phase IV MIB (decnetiv)

Charter

Chairperson:
Jon Saperia, saperia%tcpjon@decwrl.dec.com

Mailing Lists:
General Discussion: phiv-mib@jove.pa.dec.com
To Subscribe: phiv-mib-request@jove.pa.dec.com

Description of Working Group:

The DECNet Phase IV MIB Working Group will define MIB elements in the experimental portion of the MIB which correspond to standard DECNet Phase IV objects. The group will also define the access mechanisms for collecting the data and transforming it into the proper ASN.1 structures to be stored in the MIB.

In accomplishing our goals, several areas will be addressed. These include: Identification of the DECNet objects to place in the MIB, identification of the tree stucture and corresponding Object ID's for the MIB elements, Generation of the ASN.1 for these new elements, development of a proxy for non-decnet based management platforms, and a test implementation.

Goals and Milestones:

Done — Review and approve the charter and description of the working group, making any necessary changes. At that meeting, the scope of the work will be defined and individual working assignements will be made.

Sep 1991 — Review first draft document, determine necessary revisions. Follow up discussion will occur on mailing list. If possible, prototype implementation to begin after revisions have been made.

Dec 1990 — Make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail. Begin 'real' implementations.

Mar 1990 — Review final draft and if OK, give to IESG for publication as RFC.

Jul 1991	Revise document based on implementations. Ask IESG to make the revision a Draft Standard.

CURRENT MEETING REPORT

Reported by Jon Saperia/ DEC

MINUTES

A small number of us got together for the first Working Group meeting of the DECNet Phase IV MIB Working Group. At the meeting a number of items were resolved, including the charter and schedule.

1. We will need some assistance in the implementation of a proxy.
   ACTION: Steve Hunter will check to see if there is any help available at his facility.
2. There was considerable discussion about the number of DECNet Phase IV objects to support since a full implementation will have more than a hundred variables based on current estimates. The group agreed that we will consider making some objects 'optional' if the list grows too large.
   ACTION: I will attempt to produce a draft listing of objects (without the ASN.1) information within the next month or two so people can begin to review the objects.
3. We also discussed overlap with vendors who already have some DECNet MIB support in their products.
   ACTION: I am using this mailing as a request to any vendors on the list to send copies of their DECNet MIBs. I expect that most implementations will have only a very few variables implemented and they are all in the private section of their mibs so there will be no interoperability problems.
4. The next step is the actual ASN.1 encoding of the variables, given the number this is a large task. If there are any ASN.1 experts who want to help with this portion of the work, please send me mail.

ATTENDEES

Pablo Brenner	sparta!pbrenner@wnet
Stan Froyd	sfroyd@salt.acc.com
Steven Hunter	hunter@ccc.mfecc.arpa
Jonathan Saperia	saperia%tcpjon@decwrl.dec.com
Mark Sleeper	
Linda Winkler	b32357@anlvm.ctd.anl.gov

3.4.4 LAN Manager (lanman)

Charter

Chairperson:
Jim Greuel, jimg@cnd.hp.com

Mailing Lists:
General Discussion: lanmanwg@cnd.hp.com
To Subscribe: lanmanwg-request@cnd.hp.com

Description of Working Group:

This working groupis chartered to define and maintain the MIB and relevant related mechanisms needed to allow management overlap between the workgroup environment (LAN Manager based) and the enterprise environment (based on TCP/IP management).

This translates into three basic objectives:

- Define a set of management information out of the existing LAN Manager objects to allow for useful management from a TCP/IP based manager.
- Develop requirements for additional network management information, as needed, and work to extend the LAN Manager interfaces to support such information.

Goals and Milestones:

TBD	Define a set of management information out of the existing LAN Manager objects to allow for useful management from a TCP/IP based manager.
TBD	Develop requirements for additional network management information, as needed, and work to extend the LAN Manager interfaces to support such information.

CURRENT MEETING REPORT

Reported by Jim Greuel/ HP

Minutes of June 8, 1990

Lan Manager MIB I Status:

Jim Greuel summarized the current status of the first LAN Manager MIB:

An RFC decision on LM MIB I is held up due to 2 issues:

- One of the 2 LM MIB subtrees is currently specified to reside under the management object id subtree (the same one RFC 1065 – now some other RFC number I can't recall – resides in). A number of individuals within the TCP/IP network management community have problems with this.
- The IAB is concerned about vendor vs IAB control in cases where an attempt is being made to publicly define management objects for a proprietary service (e.g., LAN Manager).

The group addressed the first concern by agreeing to move all LAN Manager MIB objects into the experimental branch of the object registration tree.

Regarding the second item, Dave Crocker, the Network Management Area Director for the Internet Engineering Steering Group (IESG), informed us that the IAB and IESG are working on operating procedures for public standardization efforts that relate to proprietary objects. Ours is not the only Working Group that falls into this category, and this is being addressed as a general issue, not one solely related to us. It was concluded that, until the IAB/IESG works this out, there is not a great deal our Working Group can do except operate in as open (and visible) a manner as possible. We agreed that the LAN Manager MIB Working Group would formally submit to the IAB/IESG, through (Dave Crocker) a request that the operating guidelines/criteria for groups such as ours be defined, and that RFC status be assigned to LM MIB I as soon as possible.

Dave also pointed out that "constituency" for the Working Group, representation from multiple organizations/companies, is an important issue (though IAB/IESG has not yet determined what "adequate" constituency is). In addition, it may prove helpful to include in the working group minutes a list of companies that have stated an intent to release products based on the LM MIB. Working Group members will check if their respective companies are in a position to make such a statement.

Two minor changes to LM MIB I proposed by Dave Perkins and Evan McGinnis (in addition to the previously described object ID change) were agreed upon:

- Replace the bit field used in svSvcStatus with a table of distinct INTEGER objects. This will make it easier for the management station to interpret this data
- Remove the CMOT example. It is based on the old CMOT spec. A CMOT example can be included as a second document later if deemed necessary.

Jim Greuel, the LM MIB I editor, will submit the LM MIB I documents (with the appropriate object ID changes) to Marshall Rose for review, then to Dave for inclusion in the Internet-Draft directory.

Lan Manager MIB II

Eric Peterson of Microsoft outlined his ideas for a second LAN Manager MIB, based at least in part, on LAN Manager 2.0. He will put together an LM MIB II draft defining objects for the following areas:

- Additional file/print sharing statistics (supported by LAN Manager 1.0 as well as 2.0).
- LM 2.0 user accounting, including domain information.
- LM 2.0 fault tolerance.

We decided to use the following guidelines in defining LM MIB II:

- Define primarily read-only objects, though some writable objects will be (cautiously) considered.
- Restrict the number of objects to less than 200.

Eric will post the LM MIB II draft to the mailing list 2-3 weeks before the July 31 IETF meeting.

Next Meeting

We agreed to meet at the next IETF Meeting in Vancouver, BC on July 31 - August 3. The group will be updated on LM MIB I status and discuss the LM MIB II draft.

Attendees

Hossein Alaee	hossein_alaee@3com.com
Dave Crocker	dcrocker@nsl.dec.com
Jim Gruel	jimg%hpcndpc@hplabs.hp.com
Dwaine Kinghorn	microsoft!dwaink
Linda Kray	
Chia Chee Kuan	kuan@twg.com
Evan McGinnis	sem@bridge2.3com.com
David Perkins	dave_perkins@3com.com
Eric Peterson	
Jim Reinstedler	
Robert Rench	
Robert Ritz	
Marshall Rose	mrose@psi.com

3.4.5 Management Services Interface (msi)

Charter

Chairperson:

Oscar Newkerk, newkerk@decwet.dec.com
Sudhanshu Verma, verma@hpindbu.hp.com

Mailing Lists:

General Discussion: msiwg@decwrl.dec.com
To Subscribe: msiwg-request@decwrl.dec.com

Description of Working Group:

The objective of the Management Services Interface Working Group is to define a management services interface by which management applications may obtain access to a heterogeneous, multi-vendor, multi-protocol set of manageable objects.

The service interface is intended to support management protocols and models defined by industry and international standards bodies. As this is an Internet Engineering Task Force Working Group, the natural focus is on current and future network management protocols and models used in the Internet. However, the interface being defined is expected to be sufficiently flexible and extensible to allow support for other protocols and other classes of manageable objects. The anticipated list of protocols includes Simple Network Management Protocol (SNMP), OSI Common Management Information Protocol (CMIP), CMIP Over TCP (CMOT), Manufacturing Automation Protocol and Technical Office Protocol CMIP (MAP/TOP CMIP) and Remote Procedure Call (RPC).

Goals and Milestones:

Done	Initial version of the Internet draft placed in the Internet-Drafts directory
Done	Revised version of the draft from editing meetings placed in the Internet-Drafts directory
Aug 1990	Initial implementation of the prototype available for test.
Done	Revised draft based on the implementation experience submitted to the RFC editor.

CURRENT MEETING REPORT

Reported by Oscar Newkerk/ DEC

The MSI Working Group met to discuss editing issues with the draft API specification. The following changes were agreed to:

1. The method of passing optional arguments as an array of flags and pointers will be changed to using explicit parameters in the procedure calls. This will allow for easier implementations using RPC.
2. The MSI document will be focused on defining management application interfaces although it will include the msi_send_reply routine to allow for replies to confirmed event reports and potential manager-to-manager communications.
3. The MSI draft will be expanded to include explicit statements about the services that will be provided to supply consistent service to the management application independent of the underlying protocol. This will be provided by adding a section to the document that specifies explicit mappings from the MSI interfaces to each of the underlying protocols.

The document will be updated to reflect the results of these decisions and the new version placed in the Internet-Drafts directory.

The following are still open issues:

1. Specification of the services needed from an 'on line MIB data service'.
2. Specification of the method for supporting security in the management operations.
3. Specification of the services needed to translate agent names to addresses.
4. Bring the event section into alignment with the work in the OIM and Alertman Working Groups.

It was agreed that there would be an ad hoc MSI Working Group meeting before the IETF meeting in Vancouver to address these open issues. The meeting will be in Seattle, WA at a date to be determined by the Working Group through the mailing list.

ATTENDEES

Hossein Alaee	hossein_alaee@3com.com
Douglas Bagnall	bagnall_d@apollo.hp.com
Pablo Brenner	sparte!pbrenner@uunet
Theodore Brunner	tob@thumper.bellcore.com
Brian Handspicker	bd@vines.dec.com
Frank Kastenholtz	kasten@interlan.interlan.com
Lee LaBarre	cel@mbunix.mitre.org
Greg Minshall	minshall@kinetics.kinetics.com
Oscar Newkerk	newkerk@decwet.dec.com
David Perkins	dave_perkins@3com.com
Michael Reilly	reilly@nsl.dec.com
Jonathan Saperia	saperia%tcpjon@decwrl.dec.com
Steve Senum	sjs@network.com
Jim Sheridan	jsherida@ibm.com
Mark Sleeper	
Lou Steinberg	LOU@ARAMIS.RUTGERS.EDU
Sudhanshu Verma	verma@hpindbu.hp.com
John Vollbrecht	jrv@merit.edu
David Waitzman	djw@bbn.com

3.4.6 OSI Internet Management (oim)

Charter

Chairperson:

Lee LaBarre, cel@mbunix.mitre.org
Brian Handspicker, bd@vines.dec.com

Mailing Lists:

General Discussion: oim@mbunix.mitre.org
To Subscribe: oim-request@mbunix.mitre.org

Description of Working Group:

This working gruop will specify management information and protocols necessary to manage IP-based and OSI-based LANs and WANs in the Internet based on OSI Management standards and drafts, NIST Implementors Agreements and NMF Recommendations. It will also provide input to ANSI, ISO, NIST and NMF based on experience in the Internet, and thereby influence the final form of OSI International Standards on management.

Goals and Milestones:

TBD Develop implementors agreements for implementation of CMIP over TCP and CMIP over OSI.

TBD Develop extensions to common IETF SMI to satisfy requirements for management of the Internet using OSI management models and protocols.

TBD Develop extensions to common IETF MIB-II to satisfy requirements for management of the Internet using OSI management models and protocols.

TBD Develop prototype implementations based on protocol implementors agreements, IETF OIM Extended SMI and Extended MIB.

TBD Promote development of products based on OIM agreements.

TBD Provide input to the ANSI, ISO, NIST and NMF to influence development of OSI standards and implementors agreements.

TBD Completion of the following drafts: Implementors Agreements, Event Management, SMI Extensions, MIB Extensions, OSI Management Overview, Guidelines for the Definition of Internet Managed Objects

3.4.7 Remote LAN Monitoring (rlanmib)

Charter

Chairperson:
Mike Erlinger, mike@mti.com

Mailing Lists:
General Discussion: rlanmib@decwrl.dec.com
To Subscribe: rlanmib-request@decwrl.dec.com

Description of Working Group:

The LAN Monitoring MIB working group is chartered to define an experimental MIB for monitoring LANs.

The working group must first decide what it covers and what terminology to use. The initial thought was to investigate the characteristics of some of the currently available products (Novell's LANtern, HP's LanProbe, and Network General's Watch Dog). From this investigation MIB variables will be defined. In accomplishing our goals several areas will be addressed. These include: identification of the objects to place in the MIB, identification of the tree structure and corresponding Object ID's for the MIB elements, generation of the ASN.1 for these new elements, and a test implementation.

Goals and Milestones:

Jul 1990	Mailing list discussion of charter and collection of concerns.
Aug 1990	Discussion and final approval of charter; discussion and agreement on models and terminology. Make writing assignments.
Dec 1990	Discussion of the first draft document. Begin work on additional drafts if needed.
Mar 1990	Review latest draft of the first document and if OK give to IESG for publication as an RFC.

3.4.8 Simple Network Management Protocol (snmp)

Charter

Chairperson:

Marshall Rose, mrose@psi.com

Mailing Lists:

General Discussion: snmp-wg@nisc.nyser.net
To Subscribe: snmp-wg-request@nisc.nyser.net

Description of Working Group:

Provide a draft RFC for an enhanced backwardly compatible MIB in 4Q89 which can be implemented and interoperability tested by 1Q90 to address critical operational requirements. After multivendor testing, draft will be submitted to the RFC Editor for standardization.

Goals and Milestones:

Done	Prepare MIB-II draft
Done	Write T1 Carrier Draft
Oct 1989	Write Ethernet-Like Draft

3.4.9 Transmission Mib (transmib)

Charter

Chairperson:
John Cook, `cook@chipcom.com`

Mailing Lists:
General Discussion: `unknown`
To Subscribe: `unknown`

Description of Working Group:

The objective of the Transmission Architecture Working Group is to drive the development, documentation and testing of MIB objects for the physical and data-link layers of the OSI model. The WG attempts to consolidate redundant MIB variables from new specifications into a universal structure.

Goals and Milestones:

Ongoing	Provide a forum for vendors and users of MAC layer communications equipment.
Ongoing	Form sub-working groups of experts to define object for the following at the data-link layer: X.25, Ethernet, Token, FDDI and T1.
Done	Form a core group to evaluate the work of the sub-working groups.
Ongoing	Act as a liaison between sub-working groups and the network management protocol working groups, including SNMP, OIM, IEEE 802.1, etc.

CURRENT MEETING REPORT

The Transmission Mib working group met jointly with the FDDI Mib working group. Detailed minutes are included under the FDDI Mib entry.

3.5 OSI Integration Area

Directors: Ross Callon/DEC and Rob Hagens/University of Wisconsin

OSI Integration Area Report

The "OSI General" working group did not meet this time, largely because we didn't have any specific work to do. However, there is a good chance that we might meet next time to informally review preliminary results from the FOPG effort (see below).

The "OSI NSAP Guidelines" working group is working on producing guidelines for assignment of NSAP addresses, including determination of routing domain boundaries.

When we look at the explosive growth of network numbers in the Internet, it is clear that we need additional hierarchy in the address structure beyond the network number. Fortunately, the OSI NSAP address structure provides this additional hierarchy. This extra hierarchy is essential both to make address assignment feasible on a worldwide basis, and to allow hierarchical routing. Given that we know that we need this capability, and OSI addressing provides this capability, we are stuck with the question: "Now that we have the needed addressing capability, exactly how do we use it". The OSI addressing format allows for distribution of the authority for assigning and subdividing the address space. In many cases however, the specific administrative personnel that are assigning addresses will not know all of the technical consequences of their choice. Thus it is important to prepare and distribute guidelines which will explain the technical and administrative considerations which impact the hierarchical assignment of NSAP addresses.

There is some question as to where the guidelines should be published. once they are prepared. Possibilities include as an addition to the GOSIP specifications, as a specific "Internet" standard, or as part of the FOPG paper (see below). At this point the group is concentrating on the contents of the guidelines, with the understanding than when they are finished they will probably be fed into other efforts, rather than being propagated as an "internet standard" separate from other standardization efforts.

The "OSI-X.400" working group discussed the experimental PRMD project at the university of Wisconsin. NASA is currently supporting work at the University of Wisconsin to introduce X.400 into the Internet. This involves use of a 1984 X.400 package (ARGO) developed at the the University of Wisconsin under contract to IBM. With IBM permission, the software is being made available to university, government and non-profit sites. Over the next six months, approximately twenty sites will participate in the experiment. The Wisconsin project team will provide assistance in installing and using the software and will manage X.400 routing and name assignment. When available, the PP X.400 package, developed at University College - London, will also be incorporated into the project.

In addition, the the group discussed the use of the Internet DNS for mapping between O/R names and RFC 822 domain names. RFC 987 specifies a mechanism for mapping between X.400 (1984) O/R names and RFC 822 domain names. As described in Appendix F of RFC 987, implementatio n of these mappings requires a database which maps between X.400 O/R names and RFC 822 addresses. Assuming an asymmetric mapping, two separate relations are required, each of which maps from one format to the other.

The University of Wisconsin has modified the PP RFC 987 gateway so that it queries the DNS for mapping information. Representation of X.400 O/R addresses in RFC 987's dmn-orname notation allows storage and retrieval of O/R addresses by the Internet DNS without syntax extensions to the DNS. This mechanism is of potential use to Internet hosts acting as X.400/RFC 822 gateways.

This lead to the discussion of the possibility of using X.500 rather than DNS, which in turn lead to a discussion of the desired resurrection of the X.500 working group. It was agreed that we need to have an X.500 working group: (i) to find solutions to those aspects of directory service function which have not been fully specified in the current X.500 standard (for example: replication); (ii) to determine what should be stored in the X.500 data bases for interoperability purposes (e.g., for investigating possible use of X.500 directory services for locating transport service bridges between ISODE and pure OSI stacks, or between TP0 over X.25 and TP4 over CLNP stacks). It was pointed out that RARE already has a European X.500 group, which was doing largely overlapping work. This lead to the observation that close cooperation with the RARE group is highly desirable.

That concludes the OSI area. However, there is other work closely related to the OSI interoperability and coexistence which is not in the OSI area. One group, the "FOPG", is in fact not part of the IETF at all (and has no relationship with the IETF), but met Monday before the IETF. The FOPG (or FNC OSI Planning Group) is writing a paper outlining the various options and technical considerations for OSI transition, interoperability and coexistence, for use of a number of federal agencies making up the FNC ("Federal Networking Council"). This group made considerable progress Monday, and expects to have a rough draft of their paper in a couple of months. This draft may be distributed to the IETF-OSI mailing list before the Vancouver meeting, in order to provide an opportunity for the OSI-general working group to give informal comments on the paper to the FOPG.

3.5.1 Assignment of OSI NSAP Addresses (osinsap)

Charter

Chairperson:
Richard Colella, colella@osi3.ncsl.nist.gov

Mailing Lists:
General Discussion: ietf-osi-nsap@osi3.ncsl.nist.gov
To Subscribe: ietf-osi-nsap-request@osi3.ncsl.nist.gov

Description of Working Group:

The OSI NSAP Guidelines working group will develop guidelines for NSAP assignment and administration (aka, the care and feeding of your NSAPs).

Assuming use of existing NSAP address standards, there are two questions facing an administration:

- Do I want to be an administrative authority for allocating NSAPs?
  - how do I become an administrative authority?
    - what organizations should expect to be an "administrative authority" in the GOSIP version 2.0 address structure
    - where do I go to become an administrative authority
  - what are the administrative responsibilities involved?
    - defining and implementing assignment procedures
    - maintaining the register of NSAP assignments
    - what are the advantages/disadvantages of being an administrative authority?
- Whether NSAPS are allocated from my own or some other administrative authority, what are the technical implications of allocating the substructure of NSAPs?
  - what should be routing domains?
    - implications of being a separate routing domain (how it will affect routes, optimality of routes, firewalls and information hiding)
    - organizing routing domains by geography versus by organization versus by network topology....
  - within any routing domain, how should areas be configured?
    - (same implications as above)

Goals and Milestones:

Dec 1990 Produce a paper describing guidelines for the acquisition and administration of NSAP addresses in the Internet.

Dec 1990 Have the paper published as an RFC

Dec 1990 Have the paper incorporated, in whole or in part, into the "GOSIP Users Guide" and the FNC OSI Planning Group document

CURRENT MEETING REPORT

Reported by Philip Almquist/ Consultant

General Issues

1. Is what we do constrained by existing GSA procedures? Not much, apparently, since GSA will probably modify their planned procedures if NIST so recommends.
2. Where should the output of this group go? An RFC, most likely, and we should (and probably can) get it into the GOSIP User's Guide.
3. Have we gotten foreign comments on the draft paper? Basically, no. This may not be a problem since most foreign sites may want to get NSAP addresses from their national authorities rather than the Internet. However, comments from non-US members of the Internet are encouraged.

Discussion of the Draft Document

Unfortunately, a number of the attendees had been unable to read the draft carefully beforehand, because it was distributed in Postscript that didn't seem to be printable on some printers. The chair promised that this problem would be resolved before the next meeting.

Mary Youssef noted that the document did not adequately address how large routing domains should be or will be; nor does it discuss how interdomain routing will be accomplished. It is crucial to understand these issues if we want to design a scheme that will be practical, given the political and technical realities of the Internet. Desire for local autonomy will provide a push towards small routing domains (similar in size to IP autonomous systems), whereas the current lack of an interdomain routing protocol will provide a push towards very large routing domains (for example, a regional network and its members might form a single routing domain). Ross Callon suggested that we were overreacting to the lack of an interdomain routing protocol because Internet deployment of OSI would be slow enough that static interdomain routing would work until OSI has a real protocol for this purpose. Tony Hain disagreed, noting that DOE will have 50-100 routing domains once they deploy DECNet Phase V. Rob Hagens spoke strongly against making kludges in our design for the sake of short term workability.

Someone pointed out that, for NSAP assignment, we should be concerned about the size of administrative domains rather than routing domains. An "administrative domain" is one or more routing domains that share the same NSAP address prefix. Thus, it is the size and distribution of administrative domains that determines how large the Internet can grow before it collapses under the weight of the amount of information that must be carried around in the interdomain routing protocol. It was pointed out that the term "administrative domain" is a politically unwise choice of wording, since it suggests that members of an administrative domain have to cede administrative control of their networks to the administrative domain, when in fact the only thing that has to be centralized is the allocation of blocks of NSAP addresses.

For example, a regional network could obtain a block of NSAP addresses and become an administrative domain, allocating chunks of its block of NSAP addresses to individual campuses. The regional network would only have to advertise to backbone networks its single block of NSAP addresses (via a single prefix), rather than one or more per campus as is done in the IP world. However, there are some cases where a campus might have a good reason to use NSAP addresses that were not from the regional's block of addresses, but regionals could (and probably should) charge extra for advertising these addresses to national backbones to discourage address entropy and the resultant excessive growth of routing information in the Internet. However, we need to be sure that we don't create policies which have the side effect of making it expensive for a campus to switch WAN providers without immediately changing the NSAP addresses of all its hosts.

The discussion returned to what seems to be the central issue: information explosion. There are two approaches:

- minimize the size of the routing information that needs to be conveyed
- devote more, faster hardware to exchanging routing information

We need to find the proper balance between these two approaches. Ross Callon suggested that most sites will be Internet leaf nodes, so we probably win the most by collapsing data near the leaves of the tree. However, for sites which are very small (and there will be a lot of them) not much collapsing will be possible at the the leaf boundary, so we'll need to have further collapsing farther up the tree to get effective size reduction of the routing data about small sites.

It was pointed out that the paper uses a stylized model of the Internet (backbones, regionals, and campuses), that ignores such real world realities as back doors and mid-level networks which are not regionals (e.g., CSNET), etc. It isn't immediately clear whether the stylized model leads us to the right conclusions. Tony Hain will try to write a more realistic model.

The issue of how mobile hosts fit into an essentially geographical addressing scheme was brought up and quickly dropped because nobody had a good answer.

The issue of whether we need a temporary interdomain routing protocol for the Internet was discussed and deferred to OSI Area Directors and the OSI General Working Group. A draft version of BRP was suggested as a likely candidate.

Paul Tsuchiya presented his draft paper, "Efficient Routing Hierarchies Using Multiple Addresses". This paper describes a hierarchical address allocation scheme which very strictly mirrors the hierarchical Internet topology. Since a host's address largely determines the route used to get to it, hosts which are accessible via multiple regionals or backbones may be assigned multiple addresses, providing alternate path routing and a primitive form of policy-based routing. The group seemed to find the approach interesting but did not reach a firm conclusion about its applicability.

It was agreed that once we start to understand how to do address assignment and run OSI in the Internet we need to somehow disseminate this knowledge to the managers of at least the mid-level networks. One good way to accomplish this might be a tutorial and discussion session at a future IETF meeting.

ATTENDEES

Philip Almquist	almquist@jessica.stanford.edu
Len Bosack	bosack@mathom.cisco.com
Ross Callon	callon@erlang.dec.com
Allan Cargille	cargille@cs.wisc.edu
Martina Chan	mchan@mot.com
A. Lyman Chapin	Lyman@merit.edu
Richard Colella	colella@osi3.ncsl.nist.gov
Curtis Cox	zk0001@wnyosi4.navy.mil
Ella Gardner	epg@gateway.mitre.org
Martin Gross	martin@protolaba.dca.mil
Robert Hagens	hagens@cs.wisc.edu
Tony Hain	hain@nmfecc.arpa
David Miller	dtm@mitre.org
Cyndi Mills	cmills@bbn.com
Mark Needleman	mhnur@vccmvsa.bitnet
John Vollbrecht	jv@merit.edu
Dan Wintringham	danw@igloo.osc.edu
Richard Woundy	rwoundy@ibm.com
Mary Youssef	mary@ibm.com

3.5.2 OSI General (osigen)

Charter

Chairperson:

Rob Hagens, hagens@cs.wisc.edu
Ross Callon, callon@erlang.dec.com

Mailing Lists:

General Discussion: ietf-osi@cs.wisc.edu
To Subscribe: ietf-osi-request@cs.wisc.edu

Description of Working Group:

Help facilitate the incorporation of the OSI protocol suite into the Internet, to operate in parallel with the TCP/IP protocol suite. Facilitate the co-existence and interoperability of the TCP/IP and OSI protocol suites.

Goals and Milestones:

TBD Specify an addressing format (from those available from the OSI NSAP addressing structure) for use in the Internet. Coordinate addressing format with GOSIP version 2 and possibly other groups.

TBD Review the OSI protocol mechanisms proposed for the upcoming Berkeley release 4.4. Coordinate efforts with Berkeley.

TBD Review GOSIP. Open liaison with Government OSI Users Group (GOSIUG) for feedback of issues and concerns that we may discover.

TBD Determine what should be used short term for (i) intra-domain routing; and (ii) inter-domain routing.

TBD For interoperability between OSI end systems and TCP/IP end systems, there will need to be application layer gateways. Determine if there are any outstanding issues here.

TBD Review short term issues involved in adding OSI gateways to the Internet. Preferably, this should allow OSI and/or dual gateways to be present by the time that Berkeley release 4.4 comes out.

3.5.3 OSI-X.400 (osix400)

Charter

Chairperson:
Rob Hagens, hagens@cs.wisc.edu

Mailing Lists:
General Discussion: ietf-osi@cs.wisc.edu
To Subscribe: ietf-osi-request@cs.wisc.edu

Description of Working Group:

The IETF OSI X.400 working group is chartered to identify and provide solutions for problems encountered when operating X.400 in a dual protocol internet. This charter includes pure X.400 operational issues as well as X.400 <-> RFC 822 gateway (ala RFC 987) issues.

Goals and Milestones:

Jul 1990	Develop a scheme to alleviate the need for static RFC 987 mapping tables.

CURRENT MEETING REPORT

Reported by Rob Hagens/ University of Wisconsin

The meeting was convened by chairman Rob Hagens.

Allan Cargille (U Wisc.) briefly described the Experimental X.400 project at the University of Wisconsin. He then fielded detailed questions about the operation of the PRMD, the address structure used, etc.

Robert Hagens described a proposal that the University of Wisconsin has submitted entitled "X.400 Introduction, Support, and Planning in the Internet". This proposal has four categories. One of these categories includes support for 2 people to work on planning the strategy and operation of an Internet PRMD. The resultant plan, developed with the guidance of the X.400 Working Group, may then be adopted by the Internet community.

At the February 1990 meeting, the WG determined that a document (tentatively titled "Transition and long-term strategy for operation of X.400/MHS in the NREN" ought to be written. The funding that the University of Wisconsin seeks will provide support for 2 people to work on such a document.

Robert Hagens then described his work on use of the Domain Name System to support the storage of RFC 987/1138 address mapping tables. Use of the DNS would eliminate the need for static tables at most Internet gateway sites.

The DNS 987/1138 work will be described by an Internet Draft, to be published shortly.

ATTENDEES

Ross Callon	callon@erlang.dec.com
Allen Cargille	cargille@wisc.edu
Richard Colella	colella@osi3.ncsl.nist.gov
Curtis Cox	zk0001@wnyosi4.navy.mil
Ella Gardner	epg@gateway.mitre.org
Robert Hagens	hagens@cs.wisc.edu
Peter Kirstein	kirstein@cs.ucl.ac.uk
Marilyn Martin	martin@cdnnet.ca
Cyndi Mills	cmills@bbn.com
Mark Needleman	mhnur@uccmvsa.bitnet
Robert Shirey	shirey@mitre.org
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Cal Thixton	cthixton@next.com
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@igloo.osc.edu

3.6 Operation Area

Interim Director: Phill Gross/NRI

There are three active working groups in the Operations area – The Topology Engineering WG (TEWG), the Network Joint Monitoring WG (NJM), and the Benchmarking Methodology WG (BMWG). TEWG and NJM met in Pittsburgh, and their reports are contained in this section.

IN July 1990, BMWG submitted an Internet-Draft titled "Benchmarking Terminology", edited by Scott Bradner (Harvard University, BMWG chair). This report will be discussed at the July 31 - Aug 4 IETF meeting in Vancouver.

Both TEWG and NJM will meet in Vancouver, but with interim chairs. TEWG chair Scott Brim will be unable to attend the Vancouver meeting. Guy Almes will convene the TEWG meeting in his place. Gene Hastings will also be unable to attend the Vancouver IETF meeting (see Chair's message in these Proceedings). Phill Gross will convene a NJM meeting in Vancouver.

3.6.1 Benchmarking Methodology (bmwg)

Charter

Chairperson:
Scott Bradner, sob@harvard.harvard.edu

Mailing Lists:
General Discussion: bmwg@harvisr.harvard.edu
To Subscribe: bmwg-request@harvisr.harvard.edu

Description of Working Group:

The major goal of the Benchmark Methodology Working Group is to make a series of recommendations concerning the measurement of the performance characteristics of different classes of network equipment and software services.

Each recommendation will describe the class of equipment or service, discuss the performance characteristics that are pertinent to that class, specify a suite of performance benchmarks that test the described characteristics, as well as specify the requirements for common reporting of benchmark results.

Classes of network equipment can be broken down into two broad categories. The first deals with stand-alone network devices such as routers, bridges, repeaters, and LAN wiring concentrators. The second category includes host dependent equipment and services, such as network interfaces or TCP/IP implementations.

Once benchmarking methodologies for stand-alone devices has matured sufficiently, the group plans to focus on methodologies for testing system-wide performance, including issues such as the responsiveness of routing algorithms to topology changes.

Goals and Milestones:

Dec 1989 Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.

Feb 1989 The document will also define various classes of stand-alone network devices such as repeaters, bridges, routers, and LAN wiring concentrators as well as detail the relative importance of various performance criteria within each class.

TBD Once the community has had time to comment on the definitions of devices and performance criteria, a second document will be issued. This document will make specific recommendations regarding the suite of benchmark performance tests for each of the defined classes of network devices.

3.6.2 Topology Engineering (tewg)

Charter

Chairperson:
Scott Brim, `swb@devvax.tn.cornell.edu`

Mailing Lists:
General Discussion: `tewg@devvax.tn.cornell.edu`
To Subscribe: `tewg-request@devvax.tn.cornell.edu`

Description of Working Group:

The Topology Engineering Working Group monitors and coordinates connections between networks, particularly routing relationships.

- Monitor interconnectivity among national and international backbones and mid-level networks.
- Monitor interconnection policies with a view of moving toward a common scheme for managing interconnectivity.
- Act as a forum where network engineers and representatives of groups of networks can come together to coordinate and tune their interconnections for better efficiency of the Internet as a whole.

Goals and Milestones:

Ongoing	Reports to the Internet community will be given reflecting what we learn each quarter. This periodic report will be of use to the IETF, to FARnet, and to the CCIRN members.
Dec 1990	An immediate project is to produce an RFC which will help mid-level networks when changing their interconnectivity.

CURRENT MEETING REPORT

Reported by Scott Brim/Cornell

AGENDA

1. Report on Europe - Peter Kirstein
2. Report on the Pacific - Torben Nielsen
3. Report on RFC work - Kent England
4. Old and new issues of concern - Scott Brim
5. Internet Cartography Project - Ted Bruner

MINUTES

Coordinating international connectivity has become a significant issue for this group, first because there are few if any precedents for what is happening internationally; and second because it is all happening so fast.

Report on Europe - Peter Kirstein

Peter Kirstein described developments in Europe with special emphasis on links to North America. He pointed out the problems managing shared resource "fat pipes", multiplexed or not, where different links may have quite different use restrictions or resource allocation policies and thus, as an example, complex backup strategies. The problems here can't be solved by just a technical or just an administrative group. "Do we know how to manage bits of SPAN separated by bits of DARPA?"

Report on the Pacific - Torben Nielsen

Torben Nielsen mostly gave a status report on the Pacific. Korea is now on. At some point New Zealand will be daisy-chained to Australia and the direct link to Hawaii will be removed. Japan has multiple medium speed links; working on merging them. Taiwan soon. Australia is no longer urging Coloured Book protocols. Link to Europe within a year. Singapore, Hong Kong, Indonesia are talking; Thailand and Malaysia are interested.

Report on RFC work - Kent England

Progress is being made on the RFC for generic mid-level routing policy and "rules of thumb". New guidelines presented at the meeting were: explicitly engineer every fallback – none should be accidental; avoid routing "ties" – there should be distinct preferences, to avoid bistable situations; and the hardest problem to diagnose is oscillation. The group working on the RFC continued that night.

Old and new issues of concern - Scott Brim

We had to skip this section because we were out of time. We are continuing on the TEWG mail list.

Items that have been resolved since the last meeting:

- VMNET interaction with the Internet: VMNET is still being planned; they will now be more conscious of the physical topology of the Internet when designing their traffic flows.
- CSNet transcontinental link: Dan Long has written a routing plan which demonstrates how they are being careful not to cause routing problems with this link.
- CA*Net and its multiple connections to NSFNET (and NASA): Dennis Ferguson has written a plan describing how they will use their multiple connections.

Items that have been brought up outside the meeting so far are:

- The Army Supercomputer Network, and how it will interact with the rest of the Internet.
- Paths which have both ends in the United States but "unintentionally" travel through other countries. This same problem exists for other countries as long as they have possible fallback paths through other countries.
- NASA's ACTS satellite system and how it will interact with the Internet.

Internet Cartography Project

Worked on jointly with the NJM working group, and presented in their report.

ATTENDEES

Guy Almes	almes@rite.edu
Philip Almquist	almquist@jessica.stanford.edu
David Borman	dab@cray.com
Scott Brim	swb@devvax.tn.cornell.edu
Ted Brunner	tob@thumper.bellcore.com
David Burdelski	daveb@ftp.com
Isidrv Castineyra	isidvu@bbu.com
Tom Easterday	tem@oar.net
Kent England	kwe@bu.edu
Roger Fajman	raf@cu.nih.gov
Mark Fedor	fedor@psi.com
Dennis Ferguson	dennis@gw.ccie.utoronto.ca
Vince Fuller	fuller@jessica.stanford.edu
Jack Hahn	hahn@umds.umd.edu
Gene Hastings	hastings@psc.edu
Steve Hubert	hubert@cac.warhington.edu
Dan Jordt	danj@nwnet.net
Phil Karn	karn@thumper.bellcore.com
Peter Kirstein	Kirstein@csivcl.ac.uk
George Marshall	george@adapt.net.com
Matt Mathis	mathis@psc.edu
Don Merritt	don@brl.mil
Paul Mochapetris	pvm@151.edu
Dave O'leary	oleary@noc.sura.net
Lee Oattes	oattes@utcs.utoronto.ca
Phil Park	ppark@bbn.com
Mike Patton	map@lcs.mit.edu
Joel Repolgle	replogle@ncq.uiuc.edu
Milt Roselinsky	cmcvax!milt@hub.vcsb.edu
Karen Roubicek	roubicek@nnsc.nsf.net
Steve Storch	sstorch@bbn.com
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Kannan Varadham	kannan@oar.net
Edward Vielmetti	emv@math.lsa.umich.edu
John Vollbrecht	jrv@merit.edu
Carol Ward	cward@spot.colorado.edu
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@igloo.osc.edu
John Wobus	jmwobus@suvm.acs.syr.edu
Sze-Ying Wuu	wuu@nisc.junc.net

CA*NET /NSFNET connections

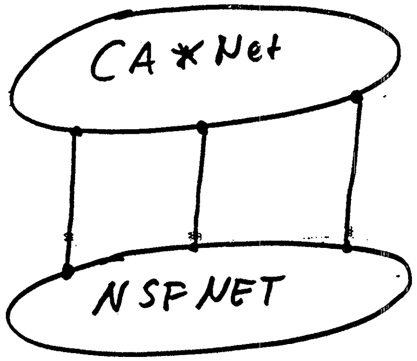

DDN/~~N~~SFNET routing

- RFC 1133
- inbound net filtering from DDN to NSFNET
- rationalization of DDN network locations
  - assign primary FIX path
  - net #10, net #26
- announcing all AS to the DDN at a high metric (20)

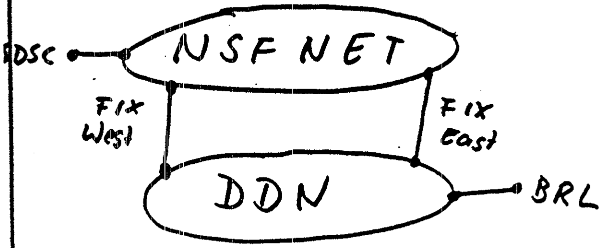

T3 demo at Net'90

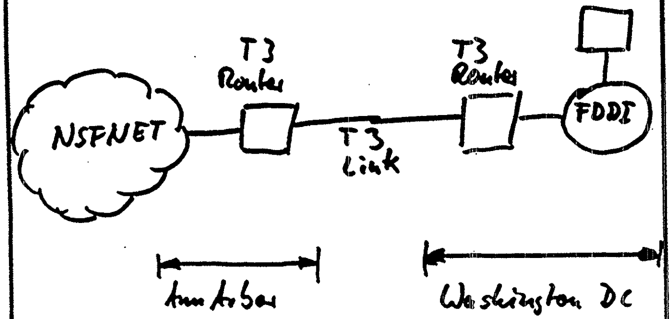

EASINET Link to CERN

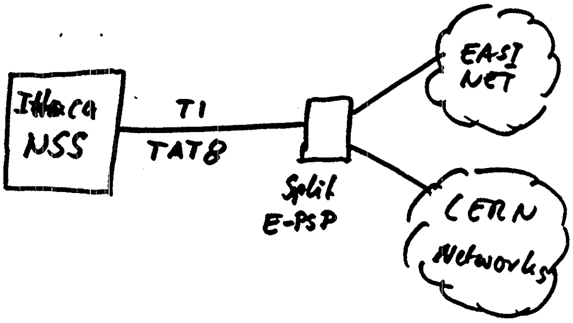

IDNX-ACSU replacements

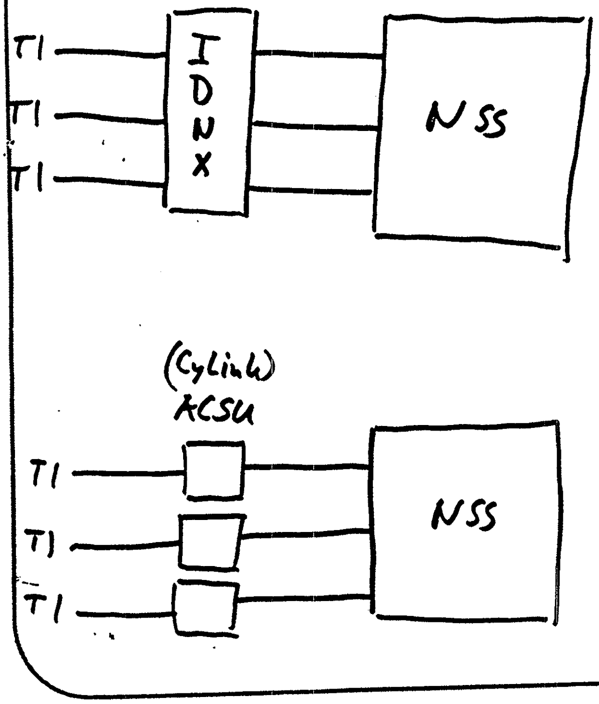

Three NSS relocations this summer.

NSS 8 Princeton, NJ
(@JvNC)

NSS 9 College Park, Md
(@U-Maryland)

NSS 10 Ithaca, NY
(@Cornell University)

Planned Third NSFNET Advanced Topics Seminar

for technical contacts at regional/peer networks

- June 21/22

3.6.3 Network Joint Management (njm)

Charter

Chairperson:
Gene Hastings, hastings@psc.edu

Mailing Lists:
General Discussion: njm@merit.edu
To Subscribe: njm-request@merit.edu

Description of Working Group:

There is a need for many different kinds of effort to deal with operational and front line engineering issues, including helping the disparate organizations work with each other. This is an attempt to solidify some of those topics. This does not make any pretense of being exhaustive.

Area of interest: operational issues and developments of the internet.

Membership: operations and engineering personnel from national backbone and mid-level networks. Other groups with responsibility for production oriented services such as security oriented groups.

Associated Technical groups: Groups which will have an interest in, and input to the agenda of this group will include the IAB and its task forces, and groups within FARnet. In particular FARnet has now several technical issues of concern, such as the selection of standard inter-network services for debugging (like maps and standard SNMP communities), and the specification of standard network statistics to be taken (of special concern is the ubiquitous ability to collect those statistics).

Meeting Times: Members of the group will represent organizations with production responsiblities. Most work will be carried on via email or teleconferencing. The group will meet at the next IETF and determine the other schedules. Sub-groups may meet between IETF meetings.

Goals and Milestones:

none specified

CURRENT MEETING REPORT

Reported by Gene Hastings/ PSC

Presentations

- Metin Feridun, BBN on Connectivity Tool (see handout) Mail to Metin Feridun `mferidun@bbn.com` if you are interested in using it. (617)-873-1870
- Ted Brunner (with Paul Tsuchiya) on cartography database proposal Propose a "MIB-ish" database format to describe router configuration and interconnection information, (see handout). `tob@thumper.bellcore.com, tsuchiya@thumper.bellcore.com`

Old Business

- DOE Community name to be announced to regional operators mailing list.
- NASA Community name to be announced to regional operators mailing list.
- Map drawing tools
  - GraphPorter, gplot (PICT>CGM>PS) GraphPorter is a commercial product for the Macintosh that converts PICT files to CGM (Computer Graphics Metafile) files. There is a companion product called MetaPICT which does the reverse. Source:

    GSC Associates
    2304 Artesia Blvd., Suite 201
    Redondo Beach, CA 90278-3114
    Phone: 213-379-2113
    Fax: 213-379-1649

    Gplot is a free collection of packages available from PSC to manipulate CGM files and render them for arbitrary output devices. Supported devices include Postscript, Tektronix, and X. Source: anonymous ftp tar file from calpe.psc.edu, in pub/gplot. Gplot does not presently run on a Macintosh, but it will run on most VMS and Unix systems. Sources are in pub/gplot/src. For inclusion in a mail distribution list, send mail to `welling@psc.edu or andrews@psc.edu` Some difficulty encountered converting MacDraw II PICT files
  - Standard statistics - still no consensus
  - efforts to cohere contact databases; Is whois up to date? Up to site administrators to make sure! Make sure there is an in-addr entry for your net, not just for hosts! -> reverse lookup for 128.182.0.0, 192.5.146.0, etc.

New Business

- Gene Hastings volunteered PSC to produce case histories of interesting or anomalous problems from their experience.

- Matt Mathis: what is happening with the phasing out of Net 10? Terrestrial Wideband showed up on a traceroute test from Cambridge Mailbridge to CMU carrying production traffic.
- Milo: Some gateways are dumb. There are no administrative routing controls within Wideband-Net.
- HWB: Recent change in NSFNET<>MILNET announcements leads to direct routes within MILNET/ARPANET which are preferred to NSFNET routes, for networks directly connected to MILNET and its close relatives, like Wideband.

It was observed that the change in NSFNET<>MILNET announcements makes sense, but it also means that network managers directly connected to MILNET or an experimental net like Wideband, need to understand this change and bias their own announcements to suit their policies.

ATTENDEES

Scott Brim	swb@devvax.tn.cornell.edu
Ted Brunner	tob@thumper.bellcore.com
Greg Dobrich	dobrich@a.psc.edu
Tom Easterday	tom@nisca.ircc.ohio-state.edu
Metin Feridun	mferidun@bbn.com
Jack Hahn	hahn@umd5.umd.edu
Gene Hastings	hastings@psc.edu
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Tom Holodnik	tjh@andrew.cmu.edu
Steven Hunter	hunter@ccc.mfecc.arpa
Robert Reschly	reschly@brl.mil
Dan Jordt	danj@cac.washington.edu
Marilyn Martin	martin@cdnnet.ca
Matt Mathis	mathis@pele.psc.edu
Milo Medin	medin@nsipo.nasa.gov
Donald Morris	morris@ucar.edu
Dave O'leary	oleary@noc.sura.net
Lee Oattes	oattes@utcs.utoronto.ca
Roxanne Streeter	streeter@nsipo.arc.nasa.gov
Paul Tsuchiya	tsuchiya@thumper.bellcore.com
Kannan Vardham	kannan@oar.net
Edward Vielmetti	emv@math.lsa.umich.edu
John Vollbrecht	jrv@merit.edu
Carol Ward	cward@spot.colorado.edu
Linda Winkler	b32357@anlvm.ctd.anl.gov
Dan Wintringham	danw@igloo.osc.edu

Sze-Ying Wuu	wuu@nisc.junc.net

Connectivity Tool

Metin Feridun
BBN STC

e-mail: mferidun@bbn.com
phone: (617) 873-1870

Connectivity Tool (CT)

- Diagnosis (location) of connectivity problems in the Internet
- Problems analyzed at the Internetwork layer (IP)
- Internet model consists of:

  hosts, gateways and networks

Current Status

- Uses a set of probe tools such as traceroute, SNMP or HMP polls, ICMP echo with IP loose source route option, etc.
- Built a platform to test diagnosis ideas for CT
  - C-based, uses object-like structures for data representation
  - X Window System based user interface (XView)
- System Requirements

  Sun 3 workstation

  SunOS 3.4, 3.5, 4.0.1 or 4.0.3

  X11R4 (or X11R3)

  ICCM compliant X Window Manager (e.g., twm)

Issues

- Protocol implementations do not always follow the standards.
  - differences across vendors and versions
  - new devices are frequently seen on the Internet
- Multiple administrative domains
  - new types of routing policies
  - frequent changes in network topology

Overall Algorithm

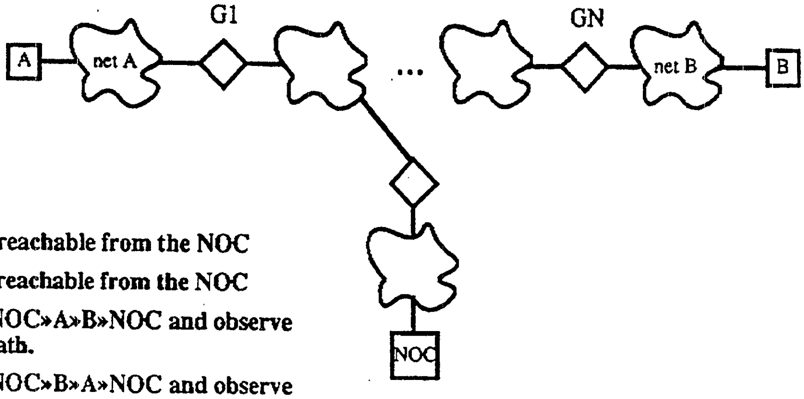

1. Verify A is reachable from the NOC
2. Verify B is reachable from the NOC
3. If (1) loop NOC»A»B»NOC and observe traversed path.
4. If (2) loop NOC»B»A»NOC and observe traversed path.

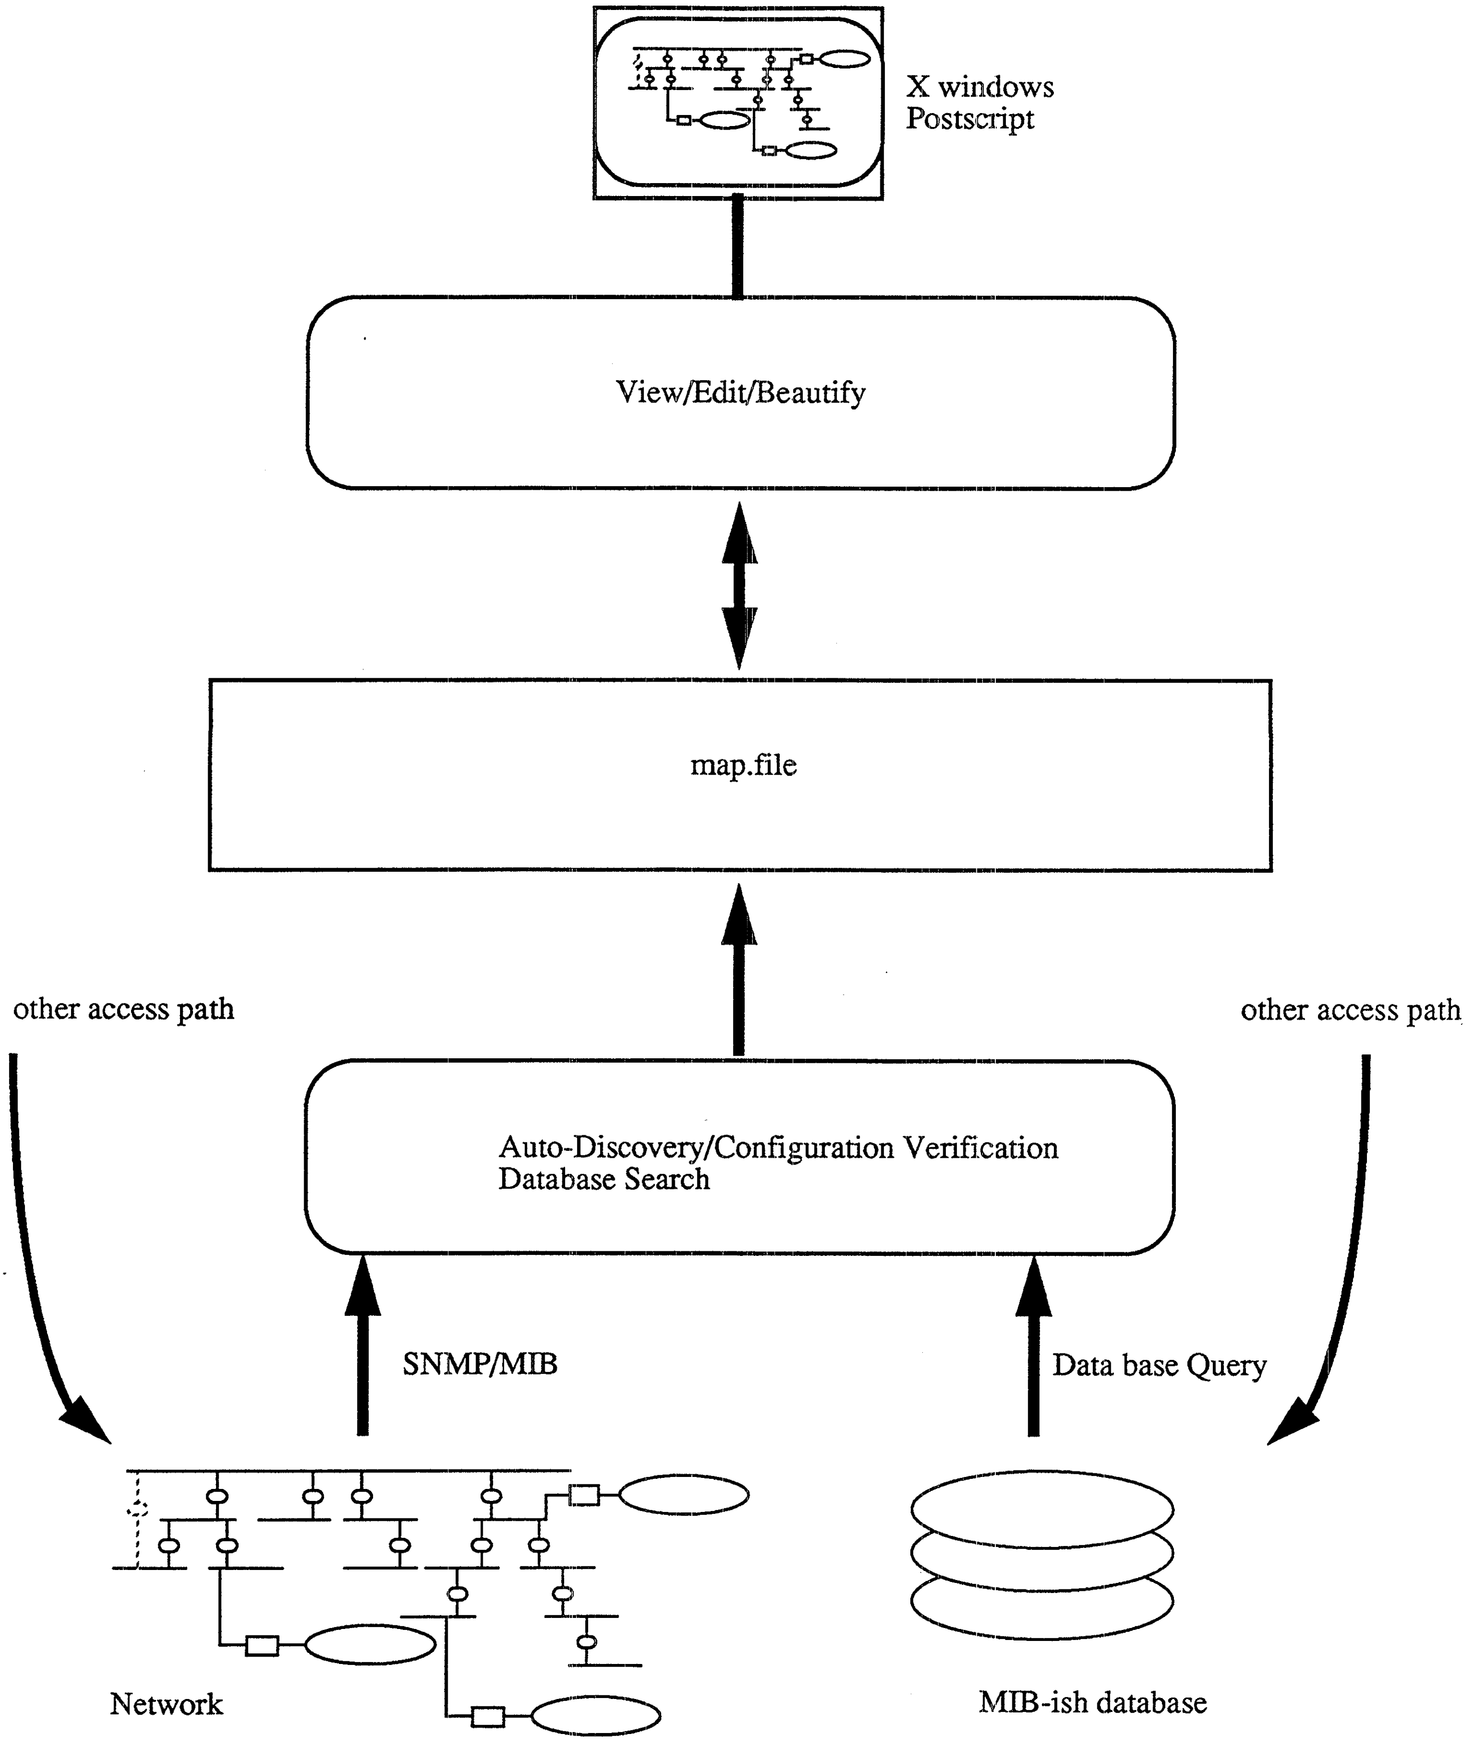
X windows
Postscript
View/Edit/Beautify
map.file
other access path
other access path
Auto-Discovery/Configuration Verification
Database Search
SNMP/MIB
Data base Query
Network
MIB-ish database

3.7 Routing Area

Director: Robert Hinden/BBN

Area Summary
Reported by Greg Vaudreuil /CNRI

The Routing area currently has five active Working Groups. Of those groups, the Interconnectivity, Multicast OSPF, and the Open Routing Working Groups held meetings in Pittsburgh. The IS-IS routing Working Group met in April and is making good progress. A revised IS-IS specification is expected to submitted by the end of the year.

The BGP protocol has been published as a proposed standard in RFC 1163 and RFC 1164. Yachov Reckter presented the latest version of the BGP protocol to the IETF. For a summary of the presentation, see the protocol presentations in section five of this document.

The Public Data Network Working Group is planning to meet again at the August IETF meeting. They intend to review the set of five documents which make of the specification. After the review, they will be submitted to be published as proposed standards.

3.7.1 ISIS for IP Internets (isis)

Charter

Chairperson:
Ross Callon, callon@erlang.dec.com

Mailing Lists:
General Discussion: isis@merit.edu
To Subscribe: isis-request@merit.edu

Description of Working Group:

The IETF IS-IS Working Group will develop additions to the existing OSI IS-IS Routing Protocol to support IP environments and dual (OSI and IP) environments.

Goals and Milestones:

Done — Develop an extension to the OSI IS-IS protocols which will allow use of IS-IS to support IP environments, and which will allow use of IS-IS as a single routing protocol to support both IP and OSI in dual environments.

TBD — Liaison with the IS-IS editor for OSI in case any minor changes to IS-IS are necessary.

TBD — Investigate the use of IS-IS to support multi-protocol routing in environments utilizing additional protocol suites.

CURRENT MEETING REPORT

Reported by Ross Callon/ DEC

Mintes of April 17th, 1990

The IS-IS working group met jointly with ANSI X3S3.3, April 17th and 18th, in Tucson Arizona. We considered two main sets of issues: (i) Possible technical changes and editorial clarifications to the IETF (integrated) IS-IS specification, based on comments received on the Internet Draft; (ii) Possible changes to the OSI IS-IS specification (for possible inclusion in the U.S. ballot comments on the ISO DP ballot).

Based on our discussions, the editor of the IETF IS-IS specification was charged with producing an updated draft, which he promised to distribute to the IS-IS working group before submission as an RFC. The main technical changes include: (i) a generalization of the manner of dealing with IP External Reachability Information in level 2 LSPs; (ii) Further definition of the Authentication field; (iii) More complete definition of the Dijkstra algorithm; and (iv) More complete definition of encapsulation and decapsulation. A number of editorial clarifications were also proposed.

In addition, our discussion resulted in one addition to the U.S. ballot comments on the ISO DP ballot. If accepted, this would allow more general treatment of external OSI reachability information, and would allow the integrated IS-IS specification to treat IP and OSI external reachable address information in the same manner.

3.7.2 Interconnectivity (iwg)

Charter

Chairperson:

Guy Almes, almes@rice.edu

Mailing Lists:

General Discussion: iwg@rice.edu
To Subscribe: iwg-request@rice.edu

Description of Working Group:

Develop the BGP protocol and BGP technical usage within the Internet, continuing the current work of the Interconnectivity Working Group in this regard.

Goals and Milestones:

Done	Complete development of version 2 of the Border Gateway Protocol (BGP).
Ongoing	Coordinate the deployment of BGP in conformance with the BGP usage document in a manner that promotes sound engineering and an open competitive environment. Take into account the interests of the various backbone and mid-level networks, the various vendors, and the user community.
Done	Develop a mature BGP technical usage document that allows us to build Inter-AS routing structures using the BGP protocol.
May 1990	Develop a MIB for BGP.
Jun 1990	Work with the Security Area to enhance the provision for security in BGP.
Jul 1990	Develop a BGP usage document describing how BGP can be used as part of a network monitoring strategy.

CURRENT MEETING REPORT

Reported by Guy Almes/ Rice

MINUTES

1. Guy Almes and Yakov Rekhter led a review of progress to date, including the conditional acceptance of the BGP Protocol document as a Proposed Internet Standard. (By mid-May, the BGP Protocol document was approved by the IESG and forwarded to the IAB for approval as a Proposed Internet Standard. Both the BGP Protocol and the BGP Usage documents will soon be published.) Changes to the protocol since the Florida State meeting were discussed.
2. Yakov Rekhter led a discussion of BGP stability. It is possible to configure a pair of neighboring ASes with incompatible routing policies such that an oscillation sets in. Yakov sketched the problem in detail and showed how the oscillation could be automatically detected.
3. Steve Willis led a discussion of a proposed MIB for BGP. This discussion resulted both in a better proposed MIB and a deeper understanding within the group of a number of BGP issues. A key issue was whether the BGP MIB should reflect the BGP information received from neighbors, actually used locally, or advertised to neighbors. Steve will follow up with an Internet Draft describing the MIB.
4. Guy Almes led a discussion of the use of BGP in monitoring the health of global Inter-AS routing. In the course of the discussion, the implications of External vs Internal BGP, even in the case of the monitoring station not being involved in routing, were shown to be important. The use of BGP for monitoring will allow a number of monitoring applications that would be totally impractical using only SNMP.
5. Guy Almes led a discussion of authentication. Consultation with members of the Security Area led to an agreement that a 16-byte Marker field per message would allow detection of spoofing. Prevention of spoofing seems to be beyond the ability of any application layered over available implementations of TCP. The presence of this 16-byte field, together with our provision of multiple authentication schemes, will allow very strong authentication. Having agreed on the need for supporting strong authentication and having modified the protocol to support it, we agreed that our needs in the near-term future were not great.

ATTENDEES

L. Allyson Brown	allyson@umd5.umd.edu
Guy Almes	almes@rice.edu
Isidro Castineyra	isidro@bbn.com
Steve Crumb	scrumb@mot.com
Robert Enger	enger@sccgate.scc.com
Dino Farinacci	dino@esd.3com.com
Dennis Ferguson	dennis@gw.ccie.utoronto.ca
Jeffrey Honig	jch@tcgould.tn.cornell.edu
Wendy Huntoon	huntoon@a.pse.edu
Peter Kirstein	kirstein@cs.ucl.ac.uk
Alex Koifman	akoifman@bbn.com
Kanchez Loa	loa@sps.mot.com
Yoni Malachi	yoni@cs.stanford.edu
C. Philip Wood	cpw@lanl.gov
Yakov Rekhter	yakov@ibm.com
Mike StJohns	stjohns@umds.umd.edu
Steve Storch	sstorch@bbn.com
Steven Willis	swillis@wellfleet.com

3.7.3 Multicast Extentions to OSPF (mospf)

Charter

Chairperson:

Steve Deering, `deering@pescadero.stanford.edu`

Mailing Lists:

General Discussion: `mospf@devvax.tn.cornell.edu`
To Subscribe: `mospf-request@devvax.tn.cornell.edu`

Description of Working Group:

This working group will extend the OSPF routing protocol so that it will be able to efficiently route IP multicast packets. This will produce a new (multicast) version of the OSPF protocol, which will be as compatible as possible with the present version (packet formats and most of the algorithms will hopefully remain unaltered).

Goals and Milestones:

Done

Become familiar with the IGMP protocol as documented in RFC 1112. Survey existing work on multicast routing, in particular, Steve Deering's paper "Multicast Routing in Internetworks and Extended LANs". Identify areas where OSPF must be extended to support multicast routing. Identify possible points of contention.

Done

Review outline of proposed changes to OSPF. Identify any unresolved issues and, if possible, resolve them.

Aug 1990

We should have a draft specification. Discuss the specification and make any necessary changes. Discuss implementation methods, using the existing BSD OSPF code, written by Rob Coltun of the University of Maryland, as an example.

Dec 1990

Report on implementations of the new multicast OSPF. Fix any problems in the specification that were found by the implementations. The specification should now be ready to submit as an RFC.

CURRENT MEETING REPORT

Reported by Steve Deering/ Stanford

Minutes:

This was the second meeting of the Multicast OSPF Working Group.

Steve Deering gave a short presentation on the cost and scaling aspects of IGMP (the Internet Group Management Protocol, specified in RFC-1112) and of link-state multicast routing as planned for OSPF. [See the accompanying slides.]

John Moy presented an outline of the changes and additions required for the current OSPF specification, to support multicast routing. Discussion of his outline identified some new concerns and suggestions regarding:

- support for multicast routing by only a subset of routers in an area and its effect on multicast vs. unicast reachability, selection of designated routers, etc.
- duplicate multicast packets arising from the support of multiple IP subnets on a single physical network.
- inter-area and inter-AS multicast routing approaches.

It was also pointed out that we will have to come up with formal MIB definitions covering the additional state and configuration variables introduced for multicast routing. No one immediately volunteered for that job.

Action Items

John Moy to generate a draft of the multicast changes/additions to the OSPF specification.

Next Meeting

The Multicast OSPF Working Group will next meet in Vancouver, at the July/August IETF meeting.

ATTENDEES

Steven Bruniges	
Duane Butler	dmb@network.com
George Clapp	meritec!clapp@bellcore.bellcore.com
Rob Coltun	rcoltun@umds.umd.edu
Steve Deering	deering@pescadero.stanford.edu
Martin Gross	gross@polaris.dca.mil
Keith Hogan	keith%penril@uunet.uu.net
Jeff Honig	jch@tcgould.tn.cornell.edu
Joseph C Lawrence	jcl@sabre.bellcore.com
Yoni Malachi	yoni@cs.stanford.edu
Tony Mason	mason@transarc.com
Jeff Mogul	mogul@decwrl.dec.com
John Moy	jmoy@proteon.com
Stephanie Price	cmcvax!price@hub.ucsb.edu
Jim Showalter	gamma@mintaka.dca.mil
Michael St. Johns	stjohns@umd5.umd.edu
Brad Strand	bstrand@cray.com
Kannan Varadhan	kannan@oar.net
Steven Willis	swillis@wellfleet.com
Chin Yuan	cyuan@srv.pacbell.com

MULTICAST OSPF

MEETING #2

Host-to-Internetwork Protocol

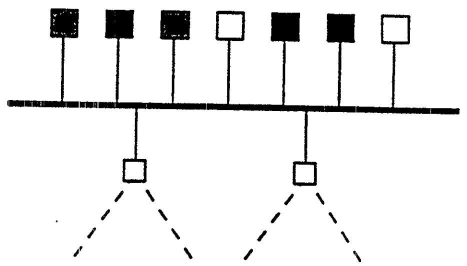

sending multicast packets:

- transmit as local multicast
- received by all members on same network
- received by attached routers, for forwarding to other networks

receiving multicast packets:

- set local address filter(s) for all groups of interest
- receive directly from senders on same network
- receive via router from senders on other networks

Host-to-Internetwork Protocol — Membership Reporting

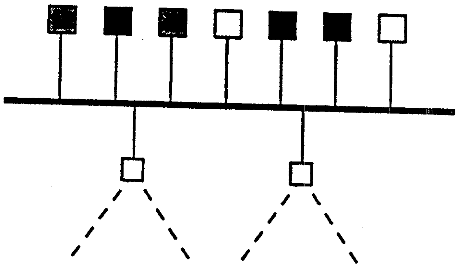

- one router per network periodically multicasts query to "all-hosts" group, scope-limited to one hop
- on receiving query, hosts set a timer for each membership to a small random interval
- when timer for group G expires, host sends a membership report *to group G*; scope-limited to one hop
- when other members of G hear report for G, they cancel their own timers.
- routers overhear all reports, and time out non-responding groups

Costs of Membership Reporting Protocol

		typical values
query interval: groups / host: groups / network:	T m M	2 minutes 5 - 20 groups 10 - 50 groups
per host		
packets processed: storage:	(1 + m) / T O(m)	3 - 10 per minute 60 - 240 bytes (IP)
per router (each interface)		
packets processed: storage:	(1 + M) / T O(M)	5 - 25 per minute 120 - 600 bytes (IP)
per network		
packets carried:	(1 + M) / T	5 - 25 per minute (< .002% Ethernet)

- negligible costs
- insensitive to size of groups

Link-State Multicast Routing

- distribute memberships with link state
- compute shortest-path tree from **source** to **member networks**, yielding:

  (source, group, in-link, out-links)

- compute on demand; cache the result
- optimization: skip if too many hops to go (out-links -> min-hops)

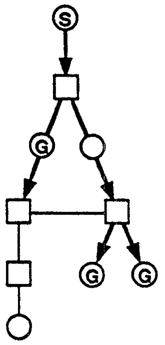

Costs of Link-State Broadcast & Multicast Routing

Cache Storage

broadcast:	~12 bytes / source seen (for IP)
multicast:	~16 bytes / (source,group) seen (for IP)

CPU Cost of a Cache Miss
(Dijkstra's Algorithm)

O(e log n)	for sparse network with e edges, n nodes
O(e)	if metric defined over a small, finite field

Example: Stanford internet
(71 networks, 48 routers, 128 edges)

Shortest Paths Computation Time

(10 mips processor)

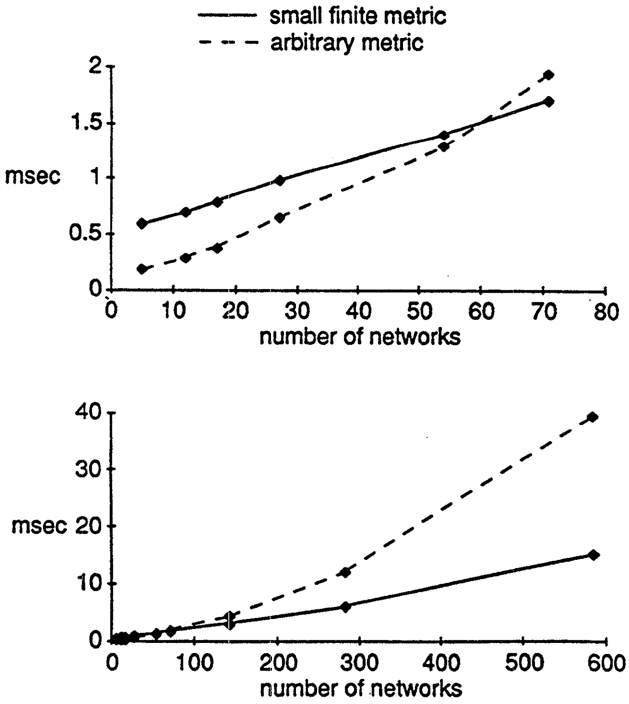

Membership Dissemination Costs for Link-State Multicast Routing

whenever

- a group appears on a network (first host joins the group)
- a group disappears from a network (last host leaves the group)

it incurs:

1 packet-hop per transit edge

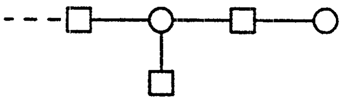

=> costs depend on rate of join/leave (per network) — difficult to predict

Membership Dissemination Costs (cont.)

Observations about join/leave rates:

- most memberships tend to be long-lived
- volatile groups tend to be sparsely distributed
- reports of disappearance can be postponed
- reports of appearance can be rate-limited

Worst-case example:

Stanford internet, 1 report / network / 5 sec

=> average 48 reports / sec
on each of 24 transit networks

(4% Ethernet @ 100 bytes / report)

Multicast OSPF spec outline

1. Conceptual Introduction
   1.1 Multicast routing fundamentals
   1.2 Characteristics
   1.3 IGMP interaction
   1.4 OSPF base
      1.4.1. Tree building
      1.4.2. Inter-area multicast
      1.4.3. Inter-AS multicast

2. Routing multicast datagrams
   2.1 Multicast cache entries

3. The multicast LSA
   3.1 When originated (IGMP)
   3.2 Function (tree pruning)
   3.3. Locality (single area only)
   3.4. Manipulation of multicast LSAs
      3.4.1 Flushing
   3.5 Wild-card multicast receivers
      3.5.1. Area border routers
      3.5.2. AS boundary routers

4. Routing cache maintenance
   4.1 Building cache entries
      4.1.1. Intra-area
         4.1.1.1 Marking routing table entries
         4.1.1.2 Tree pruning
         4.1.1.3 TTL scoping
      4.1.2 Inter-area
      4.1.3 Inter-AS
   4.2 Clearing cache entries
      4.2.1 Topological changes
         4.2.1.1 Internal
         4.2.22 External
      4.2.2 Group membership changes

5. Making multicast optional
   5.1 Router links LSA
   5.2 Flooding
      5.2.1 Indication during database exchange

6. Limitations of the protocol
   6.1 When DR not multicast-capable
   6.2 Don't require ABRs and ASBRs to be wild cards
   6.3 Single IP network/subnet per wire

Appendices

A. Configuration information
   A.1 Wild-card designation (global)
   A.2 Multicast capability (per-area)

B. Modified OSPF packet formats
   B.1 Database Description packets
   B.2 Routers links LSA

C. The multicast LSA

D. MIB

3.7.4 Open Systems Routing (orwg)

Charter

Chairperson:

Martha Steenstrup, msteenst@bbn.com

Mailing Lists:

General Discussion: open-rout-interest@bbn.com
To Subscribe: open-rout-request@bbn.com

Description of Working Group:

The Open Systems Routing Working Group is chartered to develop a policy-based AS-AS routing protocol that will accommodate large size and general topology.

Goals and Milestones:

Done	Write an architecture document.
TBD	Draft Protocol Specification of key elements of the protocol.

CURRENT MEETING REPORT

Reported by Martha Steenstrup/ BBN

During the past six months, the ORWG has generated both an architecture document and protocol specification document for inter-domain policy routing (IDPR). The former is available as an Internet Draft, and the latter will soon be submitted as an Internet Draft.

The ORWG met for three sessions at the May 1990 IETF. Our original intent was to argue about details of the IDPR protocols. However, as many new people attended our sessions, we elected to conduct the meetings as IDPR tutorials. Isidro Castineyra and Martha Steenstrup presented an overview of the architecture as well as the protocols in order to give everyone an idea of the goals and functionality of IDPR.

Prior to and during the IETF meeting, we distributed a draft of the IDPR protocol specification document and requested comments from the readers. Prototype implementation is already under way and thus the more feedback we receive now, the more likely the suggested changes will be implemented in the prototype. Please address comments to open-rout-wg@bbn.com. To be included on this list, please send a request to msteenst@bbn.com.

ATTENDEES

Doug Bagnall	bagnall_d@apollo.hp.com
Pat Barron	pat@transarc.com
Fred Bohle	fab@saturn.acc.com
Terry Braun	tab@kinetics.com
Scott Brim	swb@devvax.tn.cornell.edu
Duane Butler	dmb@network.com
Isidro Castineyra	isidro@bbn.com
Noel Chiappa	jnc@lcs.mit.edu
Dino Farinacci	dino@bridge2.3com.com
Jim Foley	jf14@vb.cc.cmu.edu
Karen Frisa	karen@kinetics.com
Olafur Gudmundsson	ocud@cs.umd.edu
Jeffrey Honig	jch@tcgould.tn.cornell.edu
Mike Horowitz	mah@shiva.com
Alex Koifman	akoifman@bbn.com
Milo Medin	medin@nsipo.nasa.gov
Donald Merritt	don@brl.mi
David Miller	dtm@mitre.org

Mark Needleman	mhnur@uccmvsa.bitnet
Phil Park	ppark@bbn.com
Yakov Rekhter	yakov@ibm.com
Don Salvin	dss@pitt.edu
Frank Solensky	solensky@interlan.interlan.com
Steve Storch	sstorch@bbn.com
Zau-Sing Su	zsu@sez.com
Ian Thomas	ian@chipcom.com
Linda Winkler	b32357@anlvm.ctd.anl.gov
C. Philip Wood	cpw@lanl.gov
Robert Woodburn	woody@saic.com
Richard Woundy	rwoundy@ibm.com
Mary Youssef	mary@ibm.com

3.7.5 Private Data Network Routing (pdnrout)

Charter

Chairperson:

CH Rokitansky, `roki@isi.edu`

Mailing Lists:

General Discussion: `pdn-wg@bbn.com`
To Subscribe: `pdn-request@bbn.com`

Description of Working Group:

The DoD INTERNET TCP/IP protocol suite has developed into a de facto industry standard for heterogenous packet switching computer networks. In the US, several hundreds of INTERNET networks are connected together; however the situation is completely different in Europe.

The only network which could be used as a backbone to allow interoperation between the many local area networks in Europe, now subscribing to the DoD INTERNET TCP/IP protocol suite, would be the system of Public Data Networks (PDN). However, so far, no algorithms have been provided to dynamically route INTERNET datagrams through X.25 public data networks. Therefore, the goals of the Public Data Network Routing working group are the development, definition and specification of required routing and gateway algorithms for an improved routing of INTERNET datagrams through the system of X.25 Public Data Networks (PDN) to allow worldwide interoperation between TCP/IP networks in various countries. In addition, the application and/or modification of the developed algorithms to interconnect local TCP/IP networks via ISDN (Integrated Services Digital Network) will be considered.

Goals and Milestones:

Done Application of the INTERNET Cluster Addressing Scheme to Public Data Networks.

Done Development of hierarchical VAN-gateway algorithms for worldwide INTERNET network reachability information exchange between VAN-gateways.

Done	Assignment of INTERNET/PDN-cluster network numbers to national public data networks. (Mapping between INTERNET network numbers and X.121 Data Network Identification Codes (DNICs)).
Done	Assignment of INTERNET/PDN-cluster addresses to PDN-hosts and VAN-gateways according to the developed hierarchical VAN-gateway algorithms.
Done	Definition of the PDN-cluster addressing scheme as an Internet standard.
Done	Specification of an X.121 Address resolution protocol.
Oct 1989	Specification of an X.25 Call Setup and Charging Determination Protocol.
Oct 1989	Specification of an X.25 Access and Forwarding Control Scheme.
Oct 1989	Specification of routing metrics taking X.25 charges into account.
TBD	Delayed TCP/IP header compression by VAN-gateways and PDN-hosts.
TBD	Provide a testbed for worldwide interoperability between local TCP/IP networks via the system of X.25 public data networks (PDN).
TBD	Implementation of the required algorithms and protocols in a VAN-Box.
TBD	Interoperability between ISO/OSI hosts on TCP/IP networks through PDN.
TBD	Consideration of INTERNET Route Servers.
TBD	Interoperability between local TCP/IP networks via ISDN.
TBD	Development of Internetwork Management Protocols for worldwide cooperation and coordination of network control and network information centers.

3.8 Security Area

Director: Steve Crocker/TIS

Area Summary

Reported by Greg Vaudreuil /CNRI

The Security Area currently has three active Working Groups. The Security Policy, SNMP Authentication and Site Security Policy handbook Working Groups met at the May IETF in Pittsburgh.

The SNMP Authentication Working Group has submitted the specification documents to the IESG for consideration as proposed standards. The protocol will be discussed at the August IETF meeting and the IESG will then forward the IETF recommendation to the IAB.

3.8.1 IP Authentication (ipauth)

Charter

Chairperson:
Jeff Schiller, jis@athena.mit.edu

Mailing Lists:
General Discussion: awg@bitsy.mit.edu
To Subscribe: awg-request@bitsy.mit.edu

Description of Working Group:

To brainstorm issues related to providing for the security and integrity of information on the Internet, with emphasis on those protocols used to operate and control the network. To propose open standard solutions to problems in network authentication.

Goals and Milestones:

TBD	RFC specifying an authentication format which supports multiple authentication systems.
TBD	Document discussing the cost/benefit tradeoffs of various generic approaches to solving the authentication problem in the Internet context.
TBD	Document to act as a protocol designers guide to authentication.
TBD	RFC proposing A Key Distribution System (emphasis on "A" as opposed to "THE"). MIT's Kerberos seems the most likely candidate here.

3.8.2 Internet Security Policy (spwg)

Charter

Chairperson:
Richard Pethia, rdp@sei.cmu.edu

Mailing Lists:
General Discussion: spwg@nri.reston.va.us
To Subscribe: spwg-request@nri.reston.va.us

Description of Working Group:

The Security Policy Working Group is chartered to create a proposed Internet Security Policy for review, possible modification, and possible adoption by the Internet Activities Board. The SPWG will focus on both technical and administrative issues related to security, including integrity, authentication and confidentiality controls, and the administration of hosts and networks.

Among the issues to be considered in this working group are:

- Responsibilities and obligations of users, data base administrators, host operators, and network managers.
- Technical controls which provide protection from disruption of service, unauthorized modification of data, unauthorized disclosure of information and unauthorized use of facilities.
- Organizational requirements for host, local network, regional network and backbone network operators.
- Incident handling procedures for various Internet components.

Goals and Milestones:

Done	Review and approve the charter making any necessary changes. Begin work on a policy framework. Assign work on detailing issues for each level of the hierarchy with first draft outline.
May 1990	Revise and approve framework documents. Begin work on detailing areas of concern, technical issues, legal issues, and recommendations for each level of the hierarchy.
Jul 1990	Prepare first draft policy recommendation for working group review and modification.

Sep 1990	Finalize draft policy and initiate review following standard RFC procedure.

CURRENT MEETING REPORT

Reported by Richard Pethia/ CERT

Minutes of the SPWG Meeting of April 17, 1990

The purpose of the April 17 meeting was to review the SPWG charter, making any necessary changes, and to begin the activity of producing a policy framework.

The initial discussion at the April 17 meeting focused on the utility of producing a security policy for the Internet, an internetwork of many networks sharing common name and address spaces. Since the Internet has no single controlling entity, and since its components are owned, operated, and administered by a variety of organizations, there was a concern that it would not be possible to enforce an Internet Security Policy in any useful way.

Despite the concerns, the attendees at this meeting decided that a formal written policy, issued by the IAB as a recommendation in the form of an RFC, could act as a vehicle to build concensus among the organizations that own and operate components of the Internet. While it was concluded that uniform policy enforcement was probably not possible, the effort of producing and promoting a security policy would benefit the Internet community by focusing attention on Internet security issues and by encouraging the component owners to take steps to improve the security of those components. In addition, the recommended policy could act as a vehicle to establish expectations of community behavior and could act as an enabling document for the development and implementation of local policy.

The group then decided that the policy should address various audiences: Internet users, host operators, network operators (including local networks, regional networks, national backbones, and international backbones), host vendors, and network vendors. For each of these audiences, the policy should speak to legal issues, technical issues, and administrative issues. Finally, the policy should, for each of the audiences, deal with the following issues: unauthorized access to data, destruction of data, modification of data, unauthorized use of service, and denial of service.

Attention then turned to the distinction between a policy and a framework to be used in developing a policy. It was generally felt that the final result of the SPWG effort should be a short, succinct document that addresses the issues listed above. The activity of developing the policy, however, should proceed using some sort of framework that would support the policy developers' efforts.

This "Internet Security Policy Development Framework" should be structured to insure all key issues are addressed and act as a working document that is elaborated over time and serves to capture the work of the policy developers.

The initial outline of the document is:

(a) Introduction
  i. Definitions and references (terms used in the balance of the document)
  ii. Internet definition
  iii. Scope of policy
  iv. Applicability
  v. Authority
  vi. Focus and emphasis

(b) Inventory of existing policies. A survey of existing policies, directives and laws that would influence an Internet Security Policy.

(c) Needed policy and architecture. A description of the audiences and issues an Internet Security Policy should address.

(d) Security Services. Covers such areas as: service classes, information classes, subscribers and users, current architectural approaches, availability, etc.

(e) Certification and Accreditation. Covers possible certification and accreditation activities including: who are the authorities, certification of components, accreditation of facilities.

(f) Security Administration and Responsibilities. Discusses issues as: overall security policy coordination, facility administration, component security administration, risk management, security training and awareness.

Minutes of the SPWG meeting of May 1, 1990

The purpose of the May 1st meeting was to discuss the policy development framework created at the April meeting and to begin work documenting areas of concern and key issues.

The framework was presented and there was general agreement that it could be used as a vehicle to develop a proposed Internet Security Policy. Discussion focused on section 4 (Security Services) of the outline and it was decided that the following three dimensions of the problem should be considered:

- Security Threats/Services
  - Confidentiality (theft of data)
  - Integrity (destruction)
  - Authentication (masquerade)
  - Assured Service (denial of service)
- Domains of Implementation
  - Administrative
  - Technical
  - Legal
- Who's Responsible

  - Users
  - Host Operators
  - Router/Network operators
  - Host Vendors
  - Router vendors

Finally, attendees brainstormed to produce the key issues listed below. Several attendees (named on individual items below) agreed to draft brief position statements on specific items in the early June time frame.

- Internet infrastructure assured service (Mike StJohns)
- User Identification - including authentication, email, remote login, ftp (Vint Cerf)
- Plugging Holes - individual responsibility (Tracy LaQuey)
- Incident Handling rules (Tracy LaQuey)
- Identification of resources (Tony Hain)
- Lines of responsibility
- User/Host/Network responsibilities (Paul Holbrook)
- Proper usage; network ethics (James Van Bokkelen)
- Configuration control
- Audit trail
- Confidentiality
- Bad Press
- User Identification - restricted access
- Denial of Service - network service
- Unauthorized access
- Adequate response when being challenged about being a source of attacks (especially when cooperating with an investigation)
- Known chain of responsible authorities
- Export restrictions - limitations enforcement

Attendees of the April Meeting

Dennis Branstad	dkb@ecf.ncsl.nist.gov
Steve Crocker	crocker@tis.com
Oma Elliott	oelliott@ddn1.dca.mil
James Ellis	ellis@psc.edu
Phill Gross	pgross@nri.reston.va.us
Paul Holbrook	ph@cert.sei.cmu.edu
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Joel Jacobs	jdj@mitre.org
Kevin Mills	mills@osi3.ncsl.nist.gov
Rich Pethia	rdp@cert.sei.cmu.edu

Rob Shirey	shirey@mitre.org
Len Tabacchi	
Greg Vaudreuil	gvaudre@nri.reston.va.us

Attendees of the May meeting

Stan Ames	sra@mbunix.mitre.org
Tom Bajzek	twb@andrew.cmu.edu
Alison Brown	alison@maverick.osc.edu
Jeffrey S. Carpenter	jjc@unix.cis.pitt.edu
Vinton Cerf	vcerf@NRI.Reston.VA.US
Richard Colella	colella@osi3.ncsl.nist.gov
Steve Crocker	crocker@tis.com
James Davin	jrd@ptt.lcs.mit.edu
Hunaid Engineer	hunaid@opus.cray.com
James Galvin	galvin@tis.com
Ella Gardner	epg@gateway.mitre.org
Tony Hain	hain@nmfecc.arpa
Robert Hoffman	hoffman@cs.pitt.edu
Paul Holbrook	ph@SEI.CMU.EDU
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Phil Karn	Karn@Thumper.Bellcore.Com
Tracy LaQuey	tracy@emx.utexas.edu
Keith McCloghrie	sytek!kzm@hplabs.hp.com
Gerald K Newman	gkn@sds.sdsc.edu
Lee Oattes	oattes@utcs.utoronto.ca
David Perkins	dave_perkins@3com.com
Marsha Perrott	mlpt@andrew.emu.edu
Richard Pethia	rdp@sei.cmu.edu
Ted Pike	tgp@sei.cmu.edu
Paul Pomes	paul_pomes@uiuc.edu
Joyce Reynolds	jkrey@venera.isi.edu
Robert J. Reschly Jr.	reschly@brl.mil
Milt Roselinsky	cmcvax!milt@hub.vcsb.edu
Jonathan Saperia	saperia%tcpjon@decwrl.dec.com
Robert W. Shirey	shirey@mitre.org
Tim Seaver	tas@mcnc.org
Michael StJohns	stjohns@umd5.umd.edu
Cal Thixton	cthixton@next.com
C. Philip Wood	cpw@lanl.gov
Sze-Ying Wuu	wuu@nisc.junc.net

3.8.3 SNMP Authentication (snmpauth)

Charter

Chairperson:
Jeff Schiller, jis@athena.mit.edu

Mailing Lists:
General Discussion: awg@bitsy.mit.edu
To Subscribe: awg-request@bitsy.mit.edu

Description of Working Group:

To define a standard mechanism for authentication within the SNMP.

Goals and Milestones:

May 1990	Write an RFC specifying procedures and formats for providing standardized authentication within the SNMP.

CURRENT MEETING REPORT

Reported by Jeff Schiller/ MIT

Mintues

The SNMP Authentication Working Group met at the Pittsburgh IETF meeting on May 2, 1990.

The primary focus of the meeting was a discussion of the relative merits of various Cryptographic Checksum algorithms used to ensure origination authentication and integrity of Protocol Data Units (PDUs). This discussion was the result of comments received from members of the Privacy and Security Research Group which reviewed the documents. Basically the problem boiled down to identifying which algorithms were both secure enough and yet were fast enough for the potential high traffic volumes that they may be needed to process. The algorithms discussed were: QMDC4, QMDC1, MD2, MD4, SNEFRU2, SNEFRU4.

It was announced at the meeting that SNEFRU2 had been broken, and the consensus was that it therefore should not be considered.

There was a sense that we needed to get cloture on the issue of what algorithm to use, in time for implementations to be demonstrated at Interop in October.

Therefore the following decisions and action items resulted:

- Consensus was reached that the RFC should *not* provide a menu of choices for implementors. Instead the RFC should specify just one of the candidate algorithms as the selected algorithm. This was argued on the basis that if more then one was allowed, each vendor would pragmatically need to support all of them, at a cost in terms of the development time for product, and memory size of the runtime binary.
- Jeff Mogul and Chuck Davin volunteered to get performance numbers on the various candidate algorithms and post their results to the mailing list. The hope here is that of all the algorithms, sufficient number would be of high performance that at least one could be found that would be both fast and secure enough to pass a review by people who can judge the security of these types of algorithms.
- The above work would be completed and a selection made in time to advance the three documents for consideration as "Proposed Standards" of the Internet.

Since the meeting was held, the performance measures have been made and it appears that MD4 is the clear performance winner. The documents will be changed to reflect this and submitted to the IETF with the recommendation they be progressed to the Proposed Draft state.

ATTENDEES

Hossein Alaee	hossein_alaee@3com.com
Stan Ames	sra@mbunix.mitre.org
Douglas Bagnall	bagnall_d@apollo.hp.com
Pat Barron	pat@trqnsarc.com
Pablo Brenner	
Alison Brown	alison@maverick@osc.edu
Ted Brunner	tob@thumper.bellcore.com
Jeff Carpenter	jjc@unix.cis.pitt.edu
Martina Chan	mchan@mot.com
Steve Crocker	crocker@tis.com
James Davin	jrd@ptt.lcs.mit.edu
Frank Kastenholtz	kasten@interlan.interlan.com
Louis Mamakos	louie@trantor.umd.edu
Keith McCloghrie	sytek!kzm@hplabs.hp.com
Jeffrey Mogul	mogul@decwrl.dec.com
Oscar Newkerk	newkerk@decwet.dec.com
John O'hara	johara@mit.edu
Brad Parker	brad@cayman.com ?
Mike Patton	map@lcs.mit.edu
David Perkins	dave_perkins@3com.com
Tod Pike	tgp@sei.emu.edu
Jonathan Saperia	saperia%tcpjon@decwrl.dec.com
Greg Satz	satz@cisco.com
Jeffrey Schiller	jis@athena.mit.edu
Richard Smith	smiddy@dss.com ?
Ted Soo-Hoo	soo-hoo@dg-rtp.dg.com
Michael StJohns	stjohns@umd5.umd.edu
Louis Steinberg	louiss@ibm.com
Ian Thomas	ian@chipcom.com
David Waiteman	djw@bbn.com
Steve Waldbusser	sw0l@andrew.cmu.edu
Y C Wang	21040 Homestead Rd Cupertino,Ca 95041

3.8.4 Site Security Policy Handbook (ssphwg)

Charter

Chairperson:

Paul Holbrook,
Joyce Reynolds, jkrey@venera.isi.edu

Mailing Lists:

General Discussion: ssphwg@cert.sei.cmu.edu
To Subscribe: ssphwg-request@cert.sei.cmu.edu

Description of Working Group:

The Site Security Policy Handbook Working Group is chartered to create a handbook that will help sites develop their own site-specific policies and procedures to deal with computer security problems and their prevention.

Among the issues to be considered in this group are:

(a) Establishing official site policy on computer security:
- Define authorized access to computing resources.
- Define what to do when local users violate the access policy.
- Define what to do when local users violate the access policy of a remote site.
- Define what to do when outsiders violate the access policy.
- Define actions to take when unauthorized activity is suspected.

(b) Establishing procedures to prevent security problems:
- System security audits.
- Account management procedures.
- Password management procedures.
- Configuration management procedures.

(c) Establishing procedures to use when unauthorized activity occurs:
- Developing lists of responsibilities and authorities: site management, system administrators, site security personnel, response teams.
- Establishing contacts with investigative agencies.
- Notification of site legal counsel.

  - Pre-defined actions on specific types of incidents (e.g., monitor activity, shut-down system).
  - Developing notification lists (who is notified of what).

(d) Establishing post-incident procedures

  - Removing vulnerabilities.
  - Capturing lessons learned.
  - Upgrading policies and procedures.

Goals and Milestones:

May 1990 Review, amend, and approve the charter as necessary. Examine the partcular customer needs for a handbook and define the scope. Continue wok on an outline for the handbook. Set up a SSPHWG "editorial board"for future writing assignments for the first draft of document.

Jun 1990 Finalize outline and organization of handbook. Partition out pieces t interested parties and SSPHWG editorial board members.

Aug 1990 Pull together a first draft handbook for working group review and modiication.

Oct 1990 Finalize draft handbook and initiate IETF Internet Draft review proces, to follow with the submission of the handbook to the RFC Editor for publication.

CURRENT MEETING REPORT

Reported by Joyce K. Reynolds/ISI and J. Paul Holbrook/ CERT

Agenda

(a) SSPHWG Charter
(b) Discussion on current security policy and relationship to the Security Policy Working Group (SPWG).
(c) Goals and directions of the SSPHWG (strawman proposal by J. Paul Holbrook)**.
**NOTE: The strawman proposal is included at the end of this report.
(d) Needs and requirements.
(e) Timeframe for writing and submission for publication of the handbook.
(f) Review of plans/action items for next round of meetings.
  i. Next meeting in Los Angeles, Tuesday, June 12th at USC/Information Sciences Institute.
  ii. Next IETF meeting in August at University of British Columbia.

Needs:

If there is a "real threat", who are the legitimate contact points:

- technical
- administrative

Phone Calls to Site(s) Three scenarios presented. You are at your site and a someone calls, stating that:

(a) They have a worm program, and would like permission to unleash it onto your site's network.
(b) They are the F.B.I., and are calling with the notification of intrusion into your site - F.B.I. suggests to keep the net open to watch the intruder.
(c) He is a hacker. The hacker states that he has capability to crack your site's passwords, etc.

What procedures and policies should be in place so that you and your site can deal with the above situations?

WHO YA GONNA CALL???

WHAT ARE THE LEGAL IMPLICATIONS??

Overview

Who are the customers of our Handbook:

- System administrators
- Site decision makers
- Site auditing the teams (?)

This Handbook will NOT be a guide on how to do:

- Risk assessment
- Contingency planning

This Handbook will promote and encourage people to hook into already existing mechanisms, even if the site doesn't have security procedures in place. They may have emergency procedures and policies that could be relevant.

Focus on things related to the network:

- Prevention
- Response
- Cleanup/followup

Assumptions:

- Network-connected
- Hosts
- Network devices
  (a) terminal servers
  (b) modems

Point out "natural" conflicts that will occur.

Physical security statement in this handbook (??) We could point out some of the risks.

- What kinds of items should be in the handbook??
- What issues should be addressed??

List and discussion of issues

(a) Physical Security

(b) Site Security Boundary
- Model : definition of terms
- Clues on what to do when you must cross organizational boundaries:
- Defining contact points
  1. technical
  2. administrative
  3. response teams
  4. investigative
- Invisible/Visible
  1. legal
  2. vendors (products or providers)
  3. press (policy and procedures)
  4. service providers

(c) Updates
- Procedures
- Tools

(d) Education of Users

(e) Historical (collection of information) [collection and protection of evidence] [facts versus assumption or —]

(f) Policy issues (Privacy)

(g) System Administrator's and Network Administrator's rights, responsibilities, AND liabilities

(h) Rights and Responsibilities of Users

(i) Formal and Informal legal procedures
- Local security/police
- FBI, Secret Service, etc.
- Verification of contact

(j) Concept of "Inter-net", "Outer-net"
- Circles of trust
- "Firewall" type concepts

(k) Procedures for working with response teams

(l) Participation in "drills" (?)

(m) "Security" of the communications lines (phones, etc.)

(n) "Insider" threats to the site

(o) Welcome banners (?)
- Is the access really authorized?
- How do you know if you're authorized?

(p) Guidelines for acceptable use?

(q) Configuration management

  - Passwords
  - Bug fixes

(r) Tools
  - Preventive
  - Response
  - Inventory of tools?

(s) Education
  - Legal
  - Administrators
  - Users (How do they deal with different kinds of threats and risks?

(t) Decision making authority
  - WHO is authorized to make what decisions?
  - WHAT authority do administrators have?
  - Layout for different cases for example: call in legal investigator, or remove a user
  - "License to hack" MUST be authorized in advance??
  - Tiger Teams

(u) Emergency response

(v) Resources
  - Other security devices
  - Other books/lists/informational sources
  - Form a subgroup?

SSPHWG volunteers will take on the task of developing a draft outline to be presented at the next SSPHWG meeting at USC/Information Sciences Institute in Marina del Rey on Tuesday, June 12th. The SSPHWG will be also be meeting at the next IETF plenary at University of British Columbia.

The following document was sent out to the SSPHWG mailing list several days before the meeting. Discussion of this document lead into the other items noted in the minutes above. There was no specific action on this document, as it was intended mainly to make sure everyone agreed with the general direction of the group.

GOALS AND DIRECTIONS OF THE SSPHWG – A STRAWMAN by J. Paul Holbrook

THE NEED

Why is there a need for a handbook like this? Looking at some of the needs may help us understand what kind of product we want to produce.

As a member of the CERT, I've come in contact with many sites trying to deal with computer security problems. It's often a rude shock when they discover someone has compromised their systems. Even for sites that have a good technical understanding of how to keep their systems secure, there are often policy questions that they haven't addressed. These policy issues make dealing with the incident much more difficult. Once the incident is over, the push to 'make sure this doesn't happen again' can result in policy made in haste; these policies can be more restrictive and cumbersome than they need to be.

A computer security compromise has much in common with any other computer 'disaster' such as an equipment breakdown or a fire. You need to have plans in place to prevent the problem, to deal with the problem while it's happening, and to deal with the consequences after the fact. Although it may seem overly-dramatic to compare a security compromise to a fire, the effect a malicious intruder may have on a site's operations could be devastating.

Another way to look at the question of 'need' is to turn it around: why should any site (especially an academic site) care about creating a computer security policy and procedures?

- There is a real threat out there. Intruders are using common holes to break into systems. Sites need to understand what the threats and risks are.
- Policies and procedures help you maintain the environment you want. Many organizations value open communication and sharing. One security incident, and "the powers that be" could force a site into a more closed environment. Policies show that you are aware of the problem and are taking steps to deal with it.
- Policies help guide cost-effective decisions. An academic site may decide that the cost of dealing with an incident doesn't warrant spending lots of time or money on defenses. A business may make a different decision.

- Policies and Procedures clarify responsibility and authority. Do you have the authority to look at a student's files? If so, do students know that?

THE CUSTOMERS OF THIS WORK

Customers of what we're trying to do:

- Systems administrators
- Site decision makers
- this includes administrators (in the traditional sense), managers, policy makers. The people who have the 'official' word on what goes on at a site

Some people who are explicitly not customers:

- Programmers
- End users

We don't need to produce a recommended set of security policies and procedures. The IETF Security Policy Working Group (SPWG) is working on that goal. Instead, whate we will produce is a set of guidelines and issues that policy makers will need to consider when they make their own policies and guidelines. This should be a tool to help people understand the need for security procedures and policies and how to go about creating them. We can include suggestions where appropriate, but much of the specifics of what a site decides to do will depend on the local circumstances. A university might make different choices about security from a government research lab.

THE OUTPUT OF THE GROUP

We hope to produce a guide to the kinds of problems sites might face, the issues they should consider, and guidelines to the kinds of steps they can take in preventing and dealing with security problems. This handbook could be published as an RFC or an FYI.

Over time, this handbook might expand to become a more general reference on site computer security. Some of the things that might be included:

- Suggested policies and procedures (perhaps whatever the Security policy WG produces)
- Bibliographies of other information to read
- Pointers to tools to help with site security

ATTENDEES

Stan Ames	sra@mbunix.mitre.org
Tom Bajzek	twb@andrew.cmu.edu
Pat Barron	pat@transarc.com
Glee Cady	ghc@merit.edu
Jeff Carpenter	jjc@unix.cis.pitt.edu
John Cavanaugh	john.cavanaugh@stpaul.ncr.com
Andrew Cherenson	arc@sgi.com
Richard Colella	colella@osi3.ncsl.nist.gov
Curtis Cox	zk0001@wnyosi4.navy.mil
Steve Crumb	scrumb@mot.com
Hunaid Engineer	hunaid@opus.cray.com
James Galvin	galvin@tis.com
Jonathan Goldick	goldick!b.psc.edu
Martyne Hallgren	martyne@tcgould.tn.cornell.edu
Greg Hollingsworth	gregh@mailer.jhuapl.edu
Tracy Laquey	tracy@emx.utexas.edu
Marilyn Martin	martin@cdnnet.ca
Donald Morris	morris@ucar.edu
Gerard Newman	gkn@sds.sdsc.edu
Marc-Andre Pepin	pepin@crim.ca
Marsha Perrott	mlp@andrew.cmu.edu
Richard Pethia	rdp@sei.cmu.edu
Tod Pike	tgp@sei.cmu.elu
Michael Reilly	reilly@nsl.dec.com
Robert Reschly	reschly@brl.mil
Tim Seaver	tas@mcnc.org
Pat Smith	psmith@merit.edu
Mary Stahl	stahl@nisc.sri.com
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
C Wood	cpw@lanl.gov
Aileen Yuan	aileen@gateway.mitre.org

Chapter 4

Network Status Briefings

4.1 State of the Internet

Presentation by Zbigniew Opalka/BBN

STATE OF THE INTERNET

Zbigniew Opalka

May 2, 1990

BBN Communications Corporation

TOPICS

- Internet Growth
- DDN Mailbridges

INTERNET GROWTH

INTERNET GROWTH SUMMARY

- 1335 Networks advertised
- 2402 Networks registered

STATE OF THE INTERNET May 2, 1990

NUMBER OF NETWORKS

December 1983-April 1990

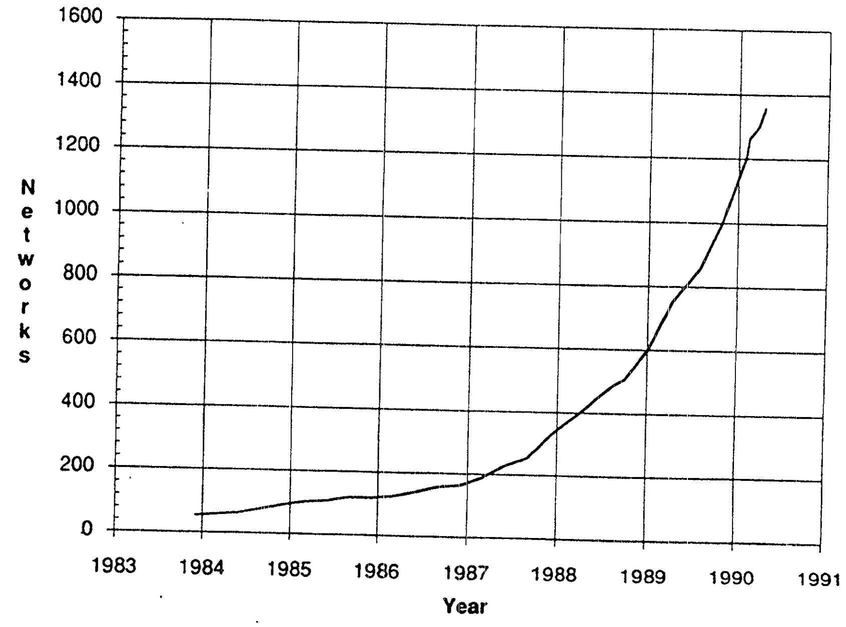

BBN Communications Corporation

STATE OF THE INTERNET May 2, 1990

NUMBER OF NETWORKS

December 1983-April 1990

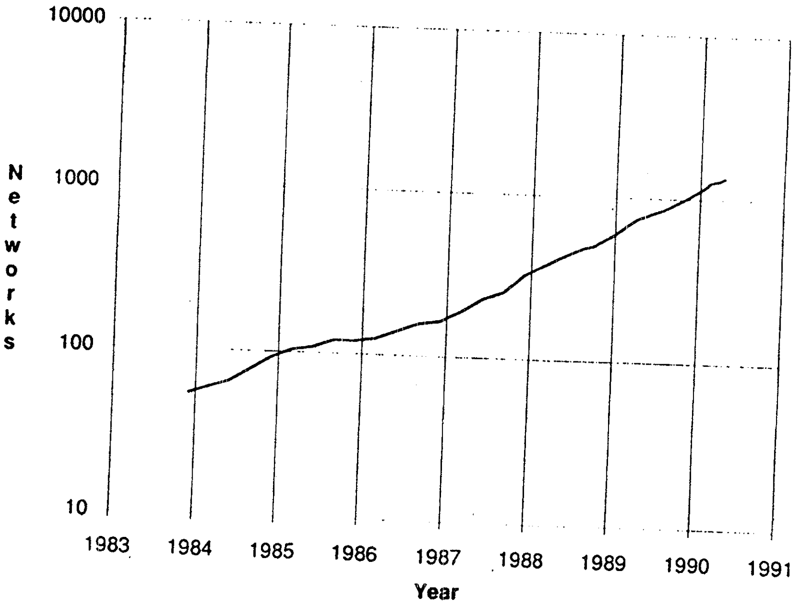

BBN Communications Corporation

STATE OF THE INTERNET May 2, 1990

ADVERTISED NETWORKS

5.0.0.0	6.0.0.0	7.0.0.0	10.0.0.0	129.83.0.0	129.85.0.0	129.87.0.0	129.88.0.0
13.0.0.0	16.0.0.0	18.0.0.0	26.0.0.0	129.89.0.0	129.91.0.0	129.92.0.0	129.93.0.0
28.0.0.0	31.0.0.0	35.0.0.0	36.0.0.0	129.95.0.0	129.96.0.0	129.97.0.0	129.99.0.0
46.0.0.0	128.2.0.0	128.3.0.0	128.4.0.0	129.100.0.0	129.101.0.0	129.104.0.0	129.105.0.0
128.5.0.0	128.6.0.0	128.7.0.0	128.8.0.0	129.106.0.0	129.107.0.0	129.108.0.0	129.109.0.0
128.9.0.0	128.10.0.0	128.11.0.0	128.16.0.0	129.110.0.0	129.111.0.0	129.112.0.0	129.114.0.0
128.18.0.0	128.19.0.0	128.20.0.0	128.29.0.0	129.116.0.0	129.117.0.0	129.118.0.0	129.119.0.0
128.32.0.0	128.35.0.0	128.36.0.0	128.38.0.0	129.120.0.0	129.121.0.0	129.122.0.0	129.123.0.0
128.39.0.0	128.41.0.0	128.42.0.0	128.43.0.0	129.124.0.0	129.125.0.0	129.126.0.0	129.127.0.0
128.44.0.0	128.46.0.0	128.47.0.0	128.49.0.0	129.128.0.0	129.129.0.0	129.130.0.0	129.132.0.0
128.52.0.0	128.54.0.0	128.55.0.0	128.56.0.0	129.133.0.0	129.137.0.0	129.138.0.0	129.139.0.0
128.58.0.0	128.59.0.0	128.60.0.0	128.61.0.0	129.140.0.0	129.142.0.0	129.151.0.0	129.152.0.0
128.62.0.0	128.63.0.0	128.81.0.0	128.82.0.0	129.170.0.0	129.171.0.0	129.172.0.0	129.173.0.0
128.83.0.0	128.84.0.0	128.86.0.0	128.89.0.0	129.174.0.0	129.175.0.0	129.176.0.0	129.177.0.0
128.91.0.0	128.92.0.0	128.93.0.0	128.95.0.0	129.179.0.0	129.186.0.0	129.188.0.0	129.189.0.0
128.96.0.0	128.97.0.0	128.99.0.0	128.100.0.0	129.190.0.0	129.191.0.0	129.192.0.0	129.194.0.0
128.101.0.0	128.102.0.0	128.103.0.0	128.104.0.0	129.195.0.0	129.199.0.0	129.200.0.0	129.206.0.0
128.105.0.0	128.109.0.0	128.110.0.0	128.111.0.0	129.207.0.0	129.213.0.0	129.216.0.0	129.217.0.0
128.112.0.0	128.113.0.0	128.114.0.0	128.115.0.0	129.218.0.0	129.219.0.0	129.221.0.0	129.229.0.0
128.116.0.0	128.117.0.0	128.118.0.0	128.119.0.0	129.231.0.0	129.235.0.0	129.236.0.0	129.237.0.0
128.120.0.0	128.121.0.0	128.122.0.0	128.123.0.0	129.238.0.0	129.240.0.0	129.241.0.0	129.243.0.0
128.124.0.0	128.125.0.0	128.126.0.0	128.127.0.0	129.244.0.0	129.245.0.0	129.246.0.0	129.248.0.0
128.129.0.0	128.134.0.0	128.135.0.0	128.136.0.0	129.250.0.0	129.252.0.0	130.11.0.0	130.13.0.0
128.138.0.0	128.139.0.0	128.140.0.0	128.141.0.0	130.14.0.0	130.15.0.0	130.17.0.0	130.18.0.0
128.143.0.0	128.144.0.0	128.145.0.0	128.146.0.0	130.20.0.0	130.33.0.0	130.34.0.0	130.37.0.0
128.147.0.0	128.148.0.0	128.149.0.0	128.150.0.0	130.38.0.0	130.39.0.0	130.41.0.0	130.42.0.0
128.151.0.0	128.152.0.0	128.153.0.0	128.154.0.0	130.44.0.0	130.46.0.0	130.49.0.0	130.50.0.0
128.155.0.0	128.156.0.0	128.157.0.0	128.158.0.0	130.53.0.0	130.54.0.0	130.56.0.0	130.57.0.0
128.159.0.0	128.160.0.0	128.162.0.0	128.163.0.0	130.58.0.0	130.59.0.0	130.60.0.0	130.62.0.0
128.164.0.0	128.165.0.0	128.167.0.0	128.168.0.0	130.63.0.0	130.68.0.0	130.69.0.0	130.70.0.0
128.169.0.0	128.170.0.0	128.171.0.0	128.172.0.0	130.74.0.0	130.82.0.0	130.84.0.0	130.85.0.0
128.173.0.0	128.174.0.0	128.175.0.0	128.178.0.0	130.86.0.0	130.87.0.0	130.90.0.0	130.91.0.0
128.180.0.0	128.181.0.0	128.182.0.0	128.183.0.0	130.92.0.0	130.93.0.0	130.94.0.0	130.95.0.0
128.184.0.0	128.185.0.0	128.186.0.0	128.187.0.0	130.102.0.0	130.105.0.0	130.108.0.0	130.109.0.0
128.188.0.0	128.189.0.0	128.190.0.0	128.192.0.0	130.111.0.0	130.113.0.0	130.114.0.0	130.116.0.0
128.193.0.0	128.194.0.0	128.195.0.0	128.196.0.0	130.118.0.0	130.123.0.0	130.125.0.0	130.126.0.0
128.197.0.0	128.198.0.0	128.200.0.0	128.202.0.0	130.127.0.0	130.132.0.0	130.134.0.0	130.137.0.0
128.204.0.0	128.205.0.0	128.206.0.0	128.208.0.0	130.144.0.0	130.150.0.0	130.152.0.0	130.153.0.0
128.209.0.0	128.210.0.0	128.211.0.0	128.212.0.0	130.155.0.0	130.157.0.0	130.160.0.0	130.163.0.0
128.213.0.0	128.214.0.0	128.215.0.0	128.217.0.0	130.166.0.0	130.167.0.0	130.168.0.0	130.180.0.0
128.218.0.0	128.219.0.0	128.220.0.0	128.221.0.0	130.182.0.0	130.184.0.0	130.185.0.0	130.188.0.0
128.222.0.0	128.223.0.0	128.226.0.0	128.227.0.0	130.189.0.0	130.191.0.0	130.194.0.0	130.195.0.0
128.228.0.0	128.229.0.0	128.230.0.0	128.231.0.0	130.199.0.0	130.201.0.0	130.202.0.0	130.203.0.0
128.235.0.0	128.236.0.0	128.237.0.0	128.238.0.0	130.204.0.0	130.207.0.0	130.208.0.0	130.212.0.0
128.239.0.0	128.241.0.0	128.242.0.0	128.244.0.0	130.215.0.0	130.217.0.0	130.219.0.0	130.220.0.0
128.245.0.0	128.247.0.0	128.248.0.0	128.249.0.0	130.221.0.0	130.223.0.0	130.224.0.0	130.225.0.0
128.250.0.0	128.252.0.0	128.253.0.0	128.255.0.0	130.232.0.0	130.233.0.0	130.235.0.0	130.236.0.0
129.1.0.0	129.3.0.0	129.5.0.0	129.6.0.0	130.237.0.0	130.238.0.0	130.239.0.0	130.240.0.0
129.7.0.0	129.8.0.0	129.10.0.0	129.13.0.0	130.241.0.0	130.245.0.0	130.252.0.0	130.253.0.0
129.15.0.0	129.16.0.0	129.18.0.0	129.19.0.0	131.1.0.0	131.2.0.0	131.4.0.0	131.5.0.0
129.20.0.0	129.21.0.0	129.22.0.0	129.24.0.0	131.6.0.0	131.7.0.0	131.8.0.0	131.9.0.0
129.25.0.0	129.26.0.0	129.29.0.0	129.30.0.0	131.10.0.0	131.11.0.0	131.12.0.0	131.13.0.0
129.32.0.0	129.33.0.0	129.34.0.0	129.40.0.0	131.14.0.0	131.15.0.0	131.16.0.0	131.17.0.0
129.43.0.0	129.46.0.0	129.47.0.0	129.48.0.0	131.18.0.0	131.19.0.0	131.20.0.0	131.21.0.0
129.49.0.0	129.51.0.0	129.53.0.0	129.55.0.0	131.22.0.0	131.23.0.0	131.24.0.0	131.25.0.0
129.57.0.0	129.59.0.0	129.60.0.0	129.61.0.0	131.26.0.0	131.27.0.0	131.28.0.0	131.29.0.0
129.62.0.0	129.63.0.0	129.64.0.0	129.65.0.0	131.30.0.0	131.31.0.0	131.32.0.0	131.33.0.0
129.66.0.0	129.69.0.0	129.71.0.0	129.72.0.0	131.34.0.0	131.35.0.0	131.36.0.0	131.37.0.0
129.73.0.0	129.74.0.0	129.75.0.0	129.77.0.0	131.38.0.0	131.39.0.0	131.40.0.0	131.41.0.0
129.78.0.0	129.79.0.0	129.81.0.0	129.82.0.0	131.42.0.0	131.43.0.0	131.44.0.0	131.45.0.0

BBN Communications Corporation

STATE OF THE INTERNET May 2, 1990

ADVERTISED NETWORKS

131.46.0.0	131.47.0.0	131.48.0.0	131.49.0.0	134.58.0.0	134.62.0.0	134.63.0.0	134.68.0.0
131.50.0.0	131.51.0.0	131.52.0.0	131.53.0.0	134.69.0.0	134.74.0.0	134.75.0.0	134.78.0.0
131.55.0.0	131.56.0.0	131.57.0.0	131.58.0.0	134.79.0.0	134.82.0.0	134.87.0.0	134.89.0.0
131.59.0.0	131.60.0.0	131.61.0.0	131.62.0.0	134.95.0.0	134.111.0.0	134.112.0.0	134.114.0.0
131.63.0.0	131.65.0.0	131.68.0.0	131.70.0.0	134.117.0.0	134.118.0.0	134.121.0.0	134.124.0.0
131.72.0.0	131.73.0.0	131.74.0.0	131.76.0.0	134.126.0.0	134.129.0.0	134.131.0.0	134.132.0.0
131.77.0.0	131.78.0.0	131.79.0.0	131.80.0.0	134.137.0.0	134.139.0.0	134.141.0.0	134.149.0.0
131.81.0.0	131.83.0.0	131.84.0.0	131.85.0.0	134.152.0.0	134.154.0.0	134.155.0.0	134.160.0.0
131.87.0.0	131.88.0.0	131.91.0.0	131.92.0.0	134.164.0.0	134.172.0.0	134.173.0.0	134.174.0.0
131.93.0.0	131.94.0.0	131.95.0.0	131.103.0.0	134.177.0.0	134.185.0.0	134.192.0.0	134.193.0.0
131.104.0.0	131.105.0.0	131.106.0.0	131.108.0.0	134.195.0.0	134.197.0.0	134.198.0.0	134.202.0.0
131.112.0.0	131.113.0.0	131.118.0.0	131.119.0.0	134.205.0.0	134.207.0.0	134.228.0.0	134.229.0.0
131.120.0.0	131.121.0.0	131.122.0.0	131.123.0.0	134.231.0.0	134.240.0.0	134.250.0.0	136.145.0.0
131.125.0.0	131.129.0.0	131.131.0.0	131.132.0.0	136.149.0.0	136.159.0.0	136.160.0.0	136.161.0.0
131.133.0.0	131.143.0.0	131.145.0.0	131.146.0.0	136.176.0.0	136.177.0.0	136.186.0.0	136.200.0.0
131.151.0.0	131.152.0.0	131.155.0.0	131.156.0.0	136.205.0.0	136.227.0.0	136.229.0.0	136.242.0.0
131.158.0.0	131.167.0.0	131.170.0.0	131.171.0.0	136.246.0.0	136.247.0.0	137.1.0.0	137.2.0.0
131.172.0.0	131.174.0.0	131.177.0.0	131.178.0.0	137.3.0.0	137.4.0.0	137.29.0.0	137.30.0.0
131.179.0.0	131.182.0.0	131.183.0.0	131.186.0.0	137.39.0.0	137.41.0.0	137.45.0.0	137.46.0.0
131.187.0.0	131.188.0.0	131.192.0.0	131.193.0.0	137.48.0.0	137.56.0.0	137.67.0.0	137.68.0.0
131.195.0.0	131.196.0.0	131.198.0.0	131.199.0.0	137.80.0.0	137.82.0.0	137.90.0.0	137.95.0.0
131.203.0.0	131.204.0.0	131.210.0.0	131.211.0.0	137.103.0.0	137.127.0.0	137.131.0.0	137.150.0.0
131.212.0.0	131.214.0.0	131.215.0.0	131.216.0.0	137.151.0.0	137.152.0.0	137.165.0.0	137.190.0.0
131.219.0.0	131.225.0.0	131.226.0.0	131.228.0.0	137.192.0.0	137.203.0.0	137.209.0.0	137.210.0.0
131.230.0.0	131.235.0.0	131.236.0.0	131.239.0.0	137.211.0.0	137.229.0.0	137.237.0.0	137.249.0.0
131.240.0.0	131.241.0.0	131.243.0.0	131.246.0.0	138.9.0.0	138.15.0.0	138.16.0.0	138.18.0.0
131.247.0.0	131.249.0.0	131.250.0.0	131.252.0.0	138.26.0.0	138.27.0.0	138.28.0.0	138.44.0.0
131.254.0.0	132.1.0.0	132.2.0.0	132.3.0.0	138.45.0.0	138.68.0.0	138.78.0.0	139.104.0.0
132.4.0.0	132.5.0.0	132.6.0.0	132.7.0.0	139.127.0.0	139.130.0.0	139.136.0.0	192.1.2.0
132.8.0.0	132.9.0.0	132.10.0.0	132.12.0.0	192.1.3.0	192.1.4.0	192.1.5.0	192.1.6.0
132.13.0.0	132.14.0.0	132.15.0.0	132.16.0.0	192.1.7.0	192.1.8.0	192.1.10.0	192.1.19.0
132.17.0.0	132.18.0.0	132.19.0.0	132.20.0.0	192.1.20.0	192.1.22.0	192.1.23.0	192.1.25.0
132.21.0.0	132.22.0.0	132.23.0.0	132.25.0.0	192.4.13.0	192.5.10.0	192.5.11.0	192.5.14.0
132.27.0.0	132.28.0.0	132.29.0.0	132.31.0.0	192.5.18.0	192.5.21.0	192.5.23.0	192.5.24.0
132.33.0.0	132.34.0.0	132.35.0.0	132.36.0.0	192.5.25.0	192.5.27.0	192.5.28.0	192.5.29.0
132.37.0.0	132.38.0.0	132.39.0.0	132.40.0.0	192.5.39.0	192.5.41.0	192.5.43.0	192.5.46.0
132.41.0.0	132.42.0.0	132.43.0.0	132.44.0.0	192.5.47.0	192.5.49.0	192.5.53.0	192.5.54.0
132.45.0.0	132.46.0.0	132.47.0.0	132.48.0.0	192.5.56.0	192.5.57.0	192.5.58.0	192.5.59.0
132.49.0.0	132.50.0.0	132.51.0.0	132.52.0.0	192.5.60.0	192.5.65.0	192.5.82.0	192.5.88.0
132.53.0.0	132.54.0.0	132.55.0.0	132.56.0.0	192.5.90.0	192.5.104.0	192.5.105.0	192.5.109.0
132.57.0.0	132.58.0.0	132.60.0.0	132.61.0.0	192.5.144.0	192.5.146.0	192.5.148.0	192.5.157.0
132.62.0.0	132.63.0.0	132.68.0.0	132.79.0.0	192.5.166.0	192.5.171.0	192.5.172.0	192.5.174.0
132.144.0.0	132.151.0.0	132.154.0.0	132.158.0.0	192.5.175.0	192.5.178.0	192.5.185.0	192.5.192.0
132.159.0.0	132.160.0.0	132.162.0.0	132.163.0.0	192.5.195.0	192.5.196.0	192.5.197.0	192.5.198.0
132.170.0.0	132.174.0.0	132.175.0.0	132.178.0.0	192.5.199.0	192.5.200.0	192.5.201.0	192.5.202.0
132.181.0.0	132.183.0.0	132.192.0.0	132.193.0.0	192.5.204.0	192.5.210.0	192.5.211.0	192.5.213.0
132.194.0.0	132.197.0.0	132.198.0.0	132.201.0.0	192.5.214.0	192.5.216.0	192.5.218.0	192.5.219.0
132.202.0.0	132.203.0.0	132.204.0.0	132.205.0.0	192.5.237.0	192.5.238.0	192.5.241.0	192.6.201.0
132.206.0.0	132.207.0.0	132.208.0.0	132.226.0.0	192.9.9.0	192.11.52.0	192.12.7.0	192.12.8.0
132.227.0.0	132.230.0.0	132.235.0.0	132.236.0.0	192.12.9.0	192.12.10.0	192.12.12.0	192.12.17.0
132.238.0.0	132.239.0.0	132.241.0.0	132.245.0.0	192.12.18.0	192.12.19.0	192.12.25.0	192.12.29.0
132.246.0.0	132.247.0.0	132.248.0.0	132.249.0.0	192.12.31.0	192.12.33.0	192.12.62.0	192.12.63.0
132.250.0.0	132.251.0.0	132.254.0.0	133.1.0.0	192.12.64.0	192.12.65.0	192.12.67.0	192.12.68.0
133.2.0.0	133.4.0.0	133.6.0.0	133.11.0.0	192.12.69.0	192.12.82.0	192.12.87.0	192.12.97.0
133.18.0.0	133.27.0.0	133.39.0.0	133.40.0.0	192.12.98.0	192.12.100.0	192.12.119.0	192.12.121.0
133.47.0.0	133.56.0.0	133.137.0.0	133.138.0.0	192.12.122.0	192.12.124.0	192.12.125.0	192.12.141.0
133.139.0.0	134.2.0.0	134.4.0.0	134.9.0.0	192.12.145.0	192.12.150.0	192.12.152.0	192.12.171.0
134.12.0.0	134.20.0.0	134.21.0.0	134.22.0.0	192.12.184.0	192.12.195.0	192.12.199.0	192.12.215.0
134.24.0.0	134.29.0.0	134.34.0.0	134.48.0.0	192.12.216.0	192.12.248.0	192.12.249.0	192.16.13.0
134.50.0.0	134.53.0.0	134.55.0.0	134.57.0.0	192.16.72.0	192.16.123.0	192.16.168.0	192.16.169.0

BBN Communications Corporation

ADVERTISED NETWORKS

192.16.172.0	192.16.173.0	192.16.174.0	192.16.175.0
192.16.184.0	192.16.188.0	192.16.201.0	192.16.202.0
192.16.204.0	192.16.205.0	192.16.207.0	192.16.209.0
192.20.225.0	192.20.239.0	192.26.8.0	192.26.10.0
192.26.12.0	192.26.18.0	192.26.20.0	192.26.25.0
192.26.26.0	192.26.27.0	192.26.49.0	192.26.83.0
192.26.86.0	192.26.88.0	192.26.91.0	192.26.93.0
192.26.147.0	192.26.148.0	192.26.200.0	192.26.210.0
192.31.3.0	192.31.7.0	192.31.8.0	192.31.17.0
192.31.21.0	192.31.24.0	192.31.27.0	192.31.30.0
192.31.39.0	192.31.44.0	192.31.63.0	192.31.66.0
192.31.69.0	192.31.70.0	192.31.71.0	192.31.75.0
192.31.82.0	192.31.83.0	192.31.85.0	192.31.87.0
192.31.93.0	192.31.95.0	192.31.100.0	192.31.103.0
192.31.104.0	192.31.106.0	192.31.111.0	192.31.112.0
192.31.146.0	192.31.147.0	192.31.153.0	192.31.165.0
192.31.172.0	192.31.179.0	192.31.181.0	192.31.192.0
192.31.197.0	192.31.211.0	192.31.215.0	192.31.222.0
192.31.223.0	192.31.225.0	192.31.230.0	192.31.231.0
192.31.239.0	192.31.242.0	192.31.253.0	192.31.254.0
192.33.4.0	192.33.5.0	192.33.6.0	192.33.9.0
192.33.13.0	192.33.14.0	192.33.19.0	192.33.33.0
192.33.36.0	192.33.112.0	192.33.115.0	192.33.116.0
192.33.140.0	192.33.141.0	192.33.144.0	192.33.145.0
192.33.146.0	192.33.148.0	192.33.149.0	192.33.153.0
192.33.156.0	192.33.159.0	192.33.167.0	192.33.168.0
192.33.170.0	192.33.178.0	192.33.179.0	192.33.181.0
192.33.182.0	192.33.185.0	192.33.212.0	192.33.213.0
192.33.214.0	192.33.215.0	192.33.216.0	192.33.217.0
192.33.219.0	192.35.44.0	192.35.49.0	192.35.59.0
192.35.62.0	192.35.74.0	192.35.75.0	192.35.76.0
192.35.78.0	192.35.79.0	192.35.82.0	192.35.86.0
192.35.100.0	192.35.129.0	192.35.140.0	192.35.142.0
192.35.147.0	192.35.148.0	192.35.154.0	192.35.162.0
192.35.163.0	192.35.165.0	192.35.167.0	192.35.169.0
192.35.170.0	192.35.171.0	192.35.180.0	192.35.196.0
192.35.200.0	192.35.208.0	192.35.213.0	192.35.226.0
192.35.229.0	192.36.23.0	192.36.125.0	192.36.148.0
192.39.11.0	192.39.12.0	192.39.14.0	192.40.51.0
192.41.140.0	192.41.146.0	192.41.177.0	192.41.192.0
192.41.197.0	192.41.200.0	192.41.202.0	192.41.204.0
192.41.211.0	192.41.217.0	192.41.228.0	192.41.237.0
192.41.245.0	192.41.246.0	192.41.249.0	192.42.2.0
192.42.4.0	192.42.41.0	192.42.61.0	192.42.62.0
192.42.66.0	192.42.70.0	192.42.80.0	192.42.81.0
192.42.82.0	192.42.88.0	192.42.91.0	192.42.95.0
192.42.108.0	192.42.110.0	192.42.114.0	192.42.141.0
192.42.142.0	192.42.144.0	192.42.152.0	192.42.153.0
192.42.155.0	192.42.239.0	192.42.244.0	192.42.245.0
192.42.246.0	192.42.247.0	192.42.248.0	192.43.152.0
192.43.188.0	192.43.197.0	192.43.199.0	192.43.204.0
192.43.205.0	192.43.207.0	192.43.217.0	192.43.239.0
192.43.244.0	192.43.250.0	192.43.252.0	192.44.1.0
192.44.82.0	192.44.83.0	192.44.84.0	192.44.85.0
192.44.216.0	192.44.217.0	192.44.220.0	192.44.221.0
192.44.222.0	192.44.223.0	192.44.224.0	192.44.225.0
192.44.233.0	192.44.234.0	192.44.235.0	192.44.236.0
192.44.237.0	192.44.253.0	192.47.242.0	192.48.33.0
192.48.80.0	192.48.96.0	192.48.100.0	192.48.114.0
192.48.115.0	192.48.118.0	192.48.125.0	192.48.139.0
192.48.143.0	192.48.153.0	192.48.212.0	192.48.214.0
192.48.215.0	192.48.217.0	192.48.219.0	192.48.224.0
192.51.37.0	192.52.61.0	192.52.64.0	192.52.65.0
192.52.66.0	192.52.71.0	192.52.106.0	192.52.107.0
192.52.111.0	192.52.117.0	192.52.154.0	192.52.156.0
192.52.159.0	192.52.160.0	192.52.179.0	192.52.180.0
192.52.182.0	192.52.194.0	192.52.195.0	192.52.218.0
192.52.219.0	192.52.220.0	192.52.232.0	192.52.233.0
192.52.236.0	192.54.33.0	192.54.81.0	192.54.104.0
192.54.106.0	192.54.109.0	192.54.138.0	192.54.148.0
192.54.222.0	192.54.226.0	192.54.240.0	192.55.87.0
192.55.90.0	192.55.97.0	192.55.98.0	192.55.103.0
192.55.106.0	192.55.109.0	192.55.117.0	192.55.120.0
192.55.133.0	192.55.134.0	192.55.187.0	192.55.188.0
192.55.190.0	192.55.207.0	192.55.225.0	192.55.228.0
192.55.229.0	192.55.234.0	192.55.235.0	192.55.245.0
192.55.246.0	192.58.1.0	192.58.2.0	192.58.3.0
192.58.91.0	192.58.107.0	192.58.109.0	192.58.127.0
192.58.150.0	192.58.151.0	192.58.152.0	192.58.153.0
192.58.181.0	192.58.194.0	192.58.199.0	192.58.204.0
192.58.206.0	192.58.219.0	192.58.222.0	192.58.223.0
192.58.224.0	192.58.225.0	192.58.232.0	192.58.233.0
192.58.234.0	192.58.235.0	192.58.244.0	192.65.50.0
192.65.71.0	192.65.78.0	192.65.97.0	192.65.137.0
192.65.141.0	192.65.143.0	192.65.147.0	192.65.175.0
192.65.176.0	192.65.177.0	192.65.185.0	192.65.202.0
192.67.6.0	192.67.53.0	192.67.62.0	192.67.67.0
192.67.70.0	192.67.80.0	192.67.92.0	192.67.97.0
192.67.128.0	192.67.129.0	192.67.134.0	192.67.175.0
192.67.176.0	192.67.177.0	192.67.184.0	192.67.225.0
192.67.226.0	192.67.227.0	192.67.228.0	192.67.250.0
192.68.26.0	192.68.132.0	192.68.160.0	192.68.162.0
192.68.163.0	192.70.131.0	192.70.132.0	192.70.138.0
192.70.187.0	192.70.236.0	192.70.246.0	

DDN MAILBRIDGES

CURRENT STATUS

- Six DDN Butterfly Mailbridges operational
- BMILAMES and BMILMTR have Ethernet interfaces to NSFNET
  - BMILAMES 192.52.195 (FIX-WEST)
  - BMILMTR 192.52.194 (FIX-EAST)
- ARPANET interface on BMILAMES, BMILISI and BMILLBL have been eliminated
- BMILISI and BMILLBL have only one interface

  - Provides EGP server function on MILNET

EGP NEIGHBORS

	DIRECT NEIGHBORS	INDIRECT NEIGHBORS
BMILAMES	90	168
BMILBBN	141	120
BMILDCEC	112	133
BMILISI	69	175
BMILLBL	43	190
BMILMTR	105	123

TRAFFIC SUMMARY

	Avg. Pkts/ Day Forwarded	Avg. Bytes/ Pkt.	Avg. Pkt. Dropped
BMILAMES	4,460,790	144	2.1%
BMILBBN	2,539,730	131	3.2%
BMILDCEC	2,648,190	138	2.7%
BMILISI	1,552,510	227	0.1%
BMILLBL	224,139	397	0.0%
BMILMTR	3,581,250	149	0.9%

BMILAMES DAILY THROUGHPUT

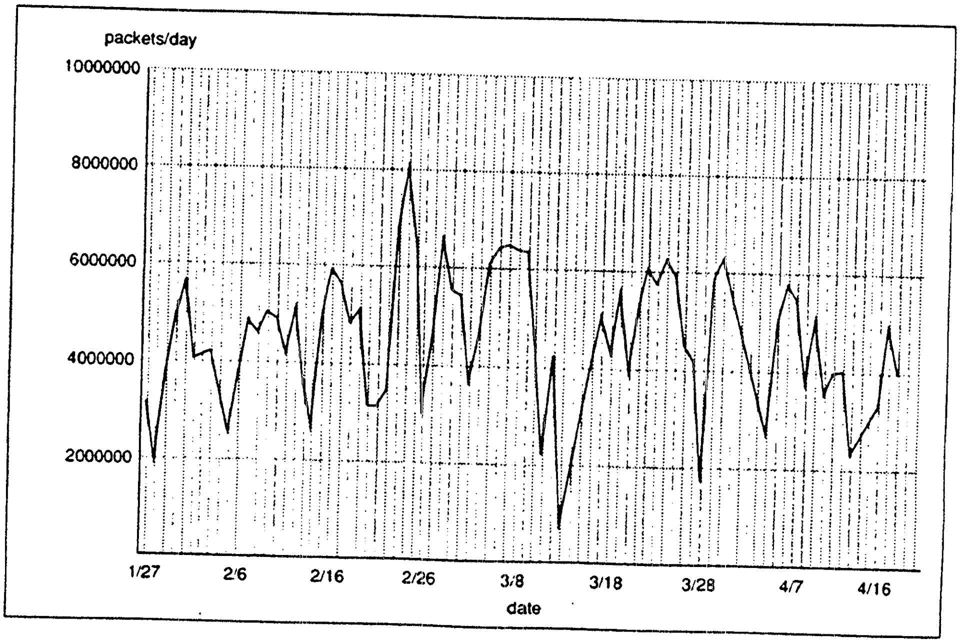

BMILBBN DAILY THROUGHPUT

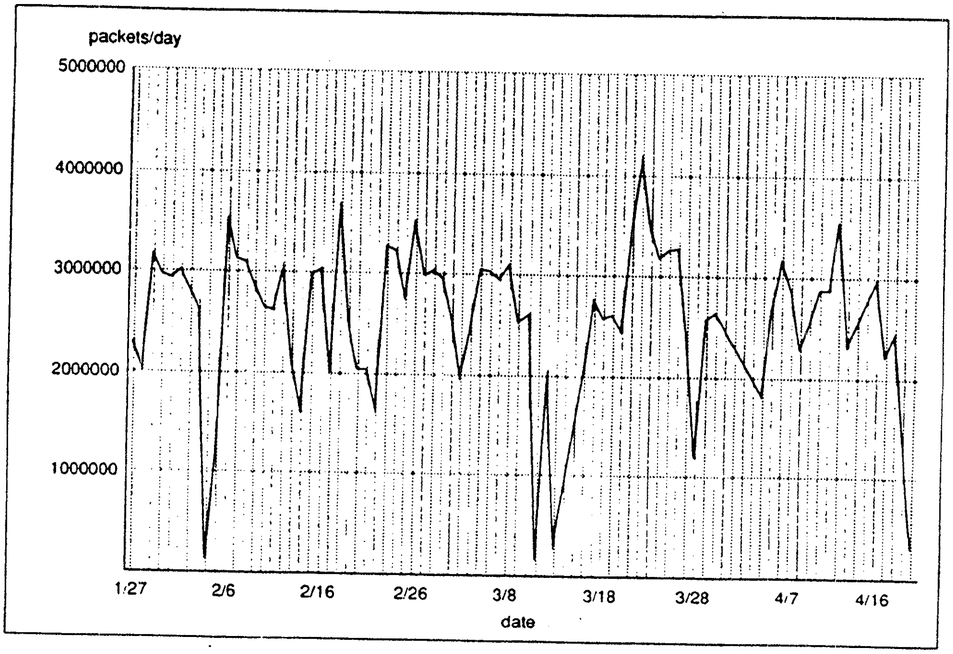

BMILDCEC DAILY THROUGHPUT

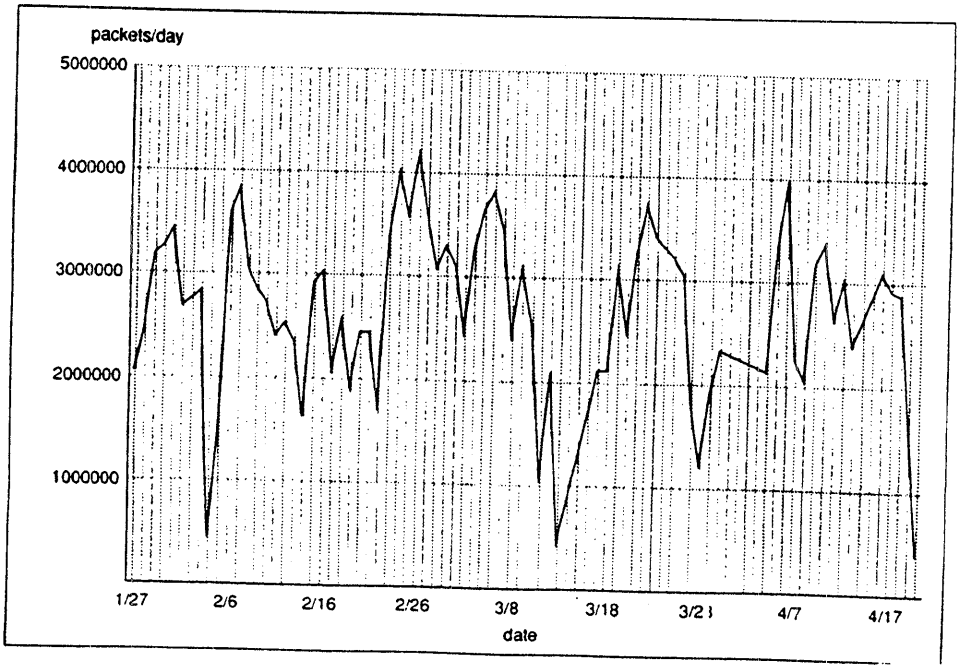

BMILISI DAILY THROUGHPUT

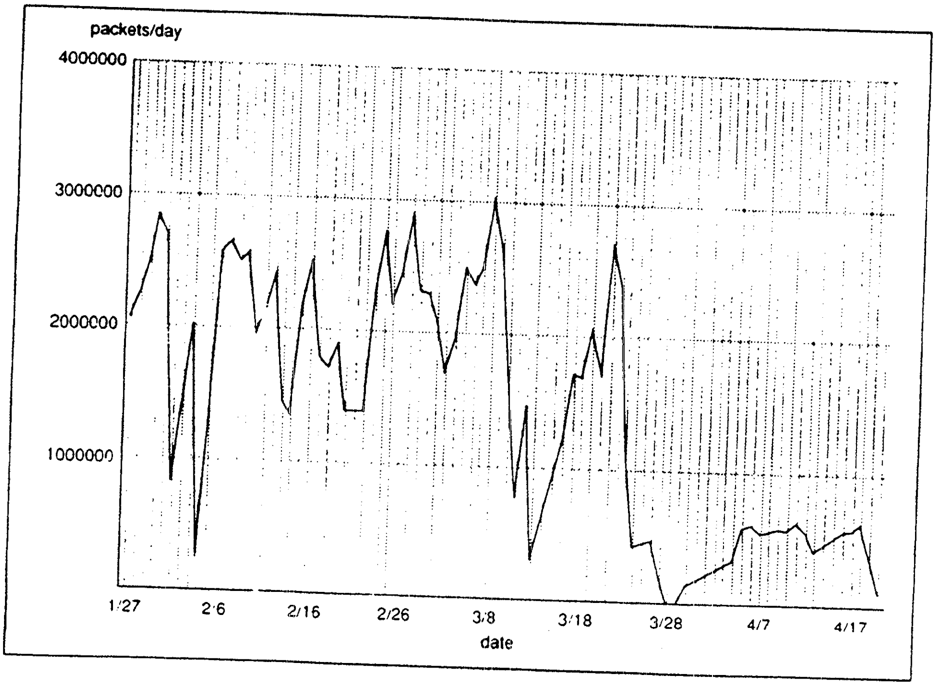

BMILLBL DAILY THROUGHPUT

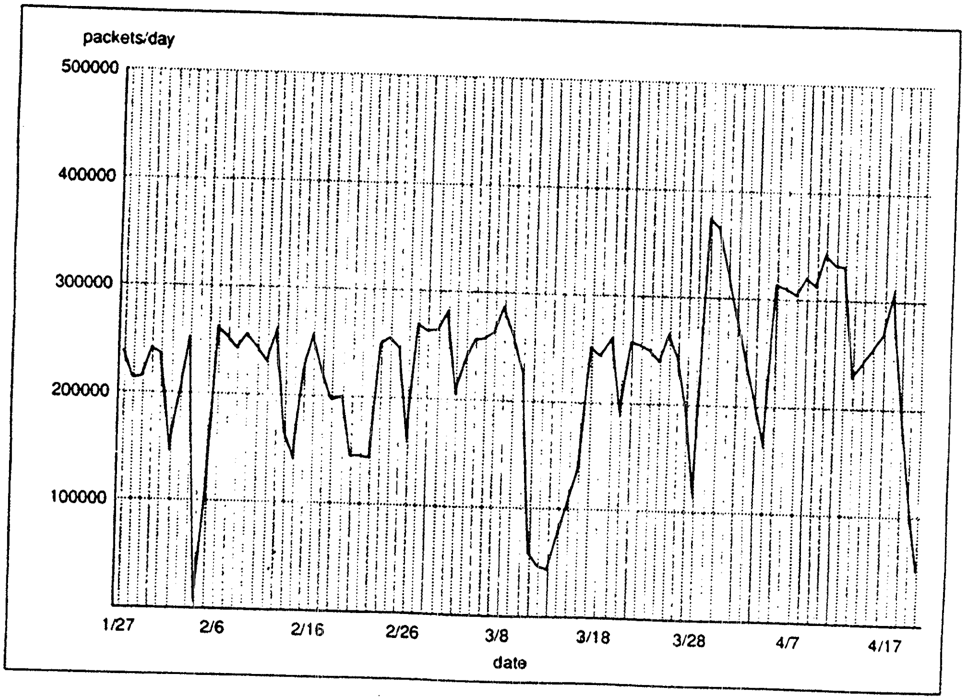

BMILMTR DAILY THROUGHPUT

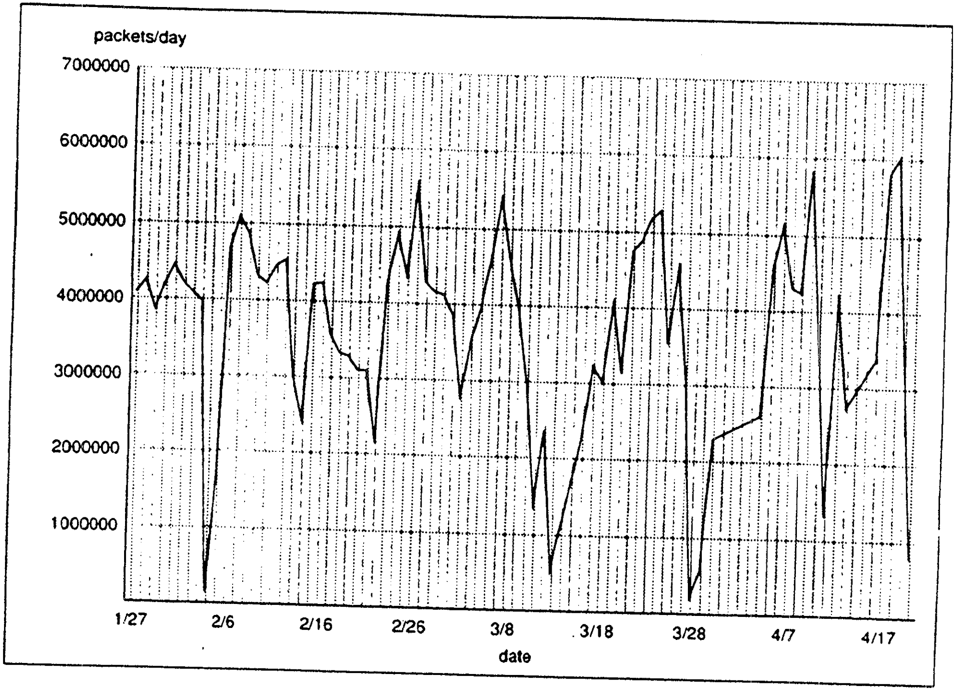

PROBLEMS AND ISSUES

- Routing difficulties (net unreachable, routing loops)
  - Dismantling of ARPANET
  - Rapid expansion of INTERNET
  - Loss of routing update due to too many EGP neighbors (> 90)
  - Unbalanced EGP neighbors distribution
- BBN and DCA are working diligently to resolve these problems
  - Deployment of the seventh Mailbridge
  - Distribute updated EGP assignment list
  - Improve Mailbridge EGP processing performance
  - Increase polling for update interval for AF concentrators

4.2 Energy Sciences Network Report

Presentation by Tony Hain/ ESnet

At the last IETF we had just received the 16 FTS 2000 circuits from AT&T. After a short shakedown, the circuits have proven stable and routine outages have been handled. We spent a fair amount of time in joint meetings with our sites and their connected regional networks to resolve the details of implementing our routing plan. On the 19th of April we turned on advertising between the sites and several regionals. We had already established peering with NSF, NSN, and the Mail Bridge at NASA Ames and have since established a peering point at UMD/SURAnet. The regionals we are currently peering with include, SURAnet, BARRnet, THEnet, SESQUInet, CERFnet and Los Nettos. We plan to peer with additional regionals as new circuits are installed over the next several months. These include, CICnet, NEARnet, NWnet, WESTnet, MIDnet and JVNCnet. We will also be installing additional sites at DOE-HQ Germantown, Md., Sandia National Lab (Abq. & Liv.), SAIC and AMES Lab Iowa. Also over the coming months we will be upgrading our cisco routers to the new CSC-3 processors.

We have seen a growth in both packet and byte counts over the first 3 months. This was likely "testing-the-waters" traffic as the sites were required to establish bidirectional static routes to our routers until April 19. The traffic distribution at this point looks to be about evenly divided between IP and DECnet, but the coming month will establish a true production baseline.

In the area of tools, Alan Sturtevant completed a DECnet trace-route which works through both DEC and cisco routers. This provides 3rd party point-to-point, circuit name and cost at each hop, and return path information. At this time the cisco support relies on telnet, but it will be moved to SNMP as soon as the required data is available in the MIB.

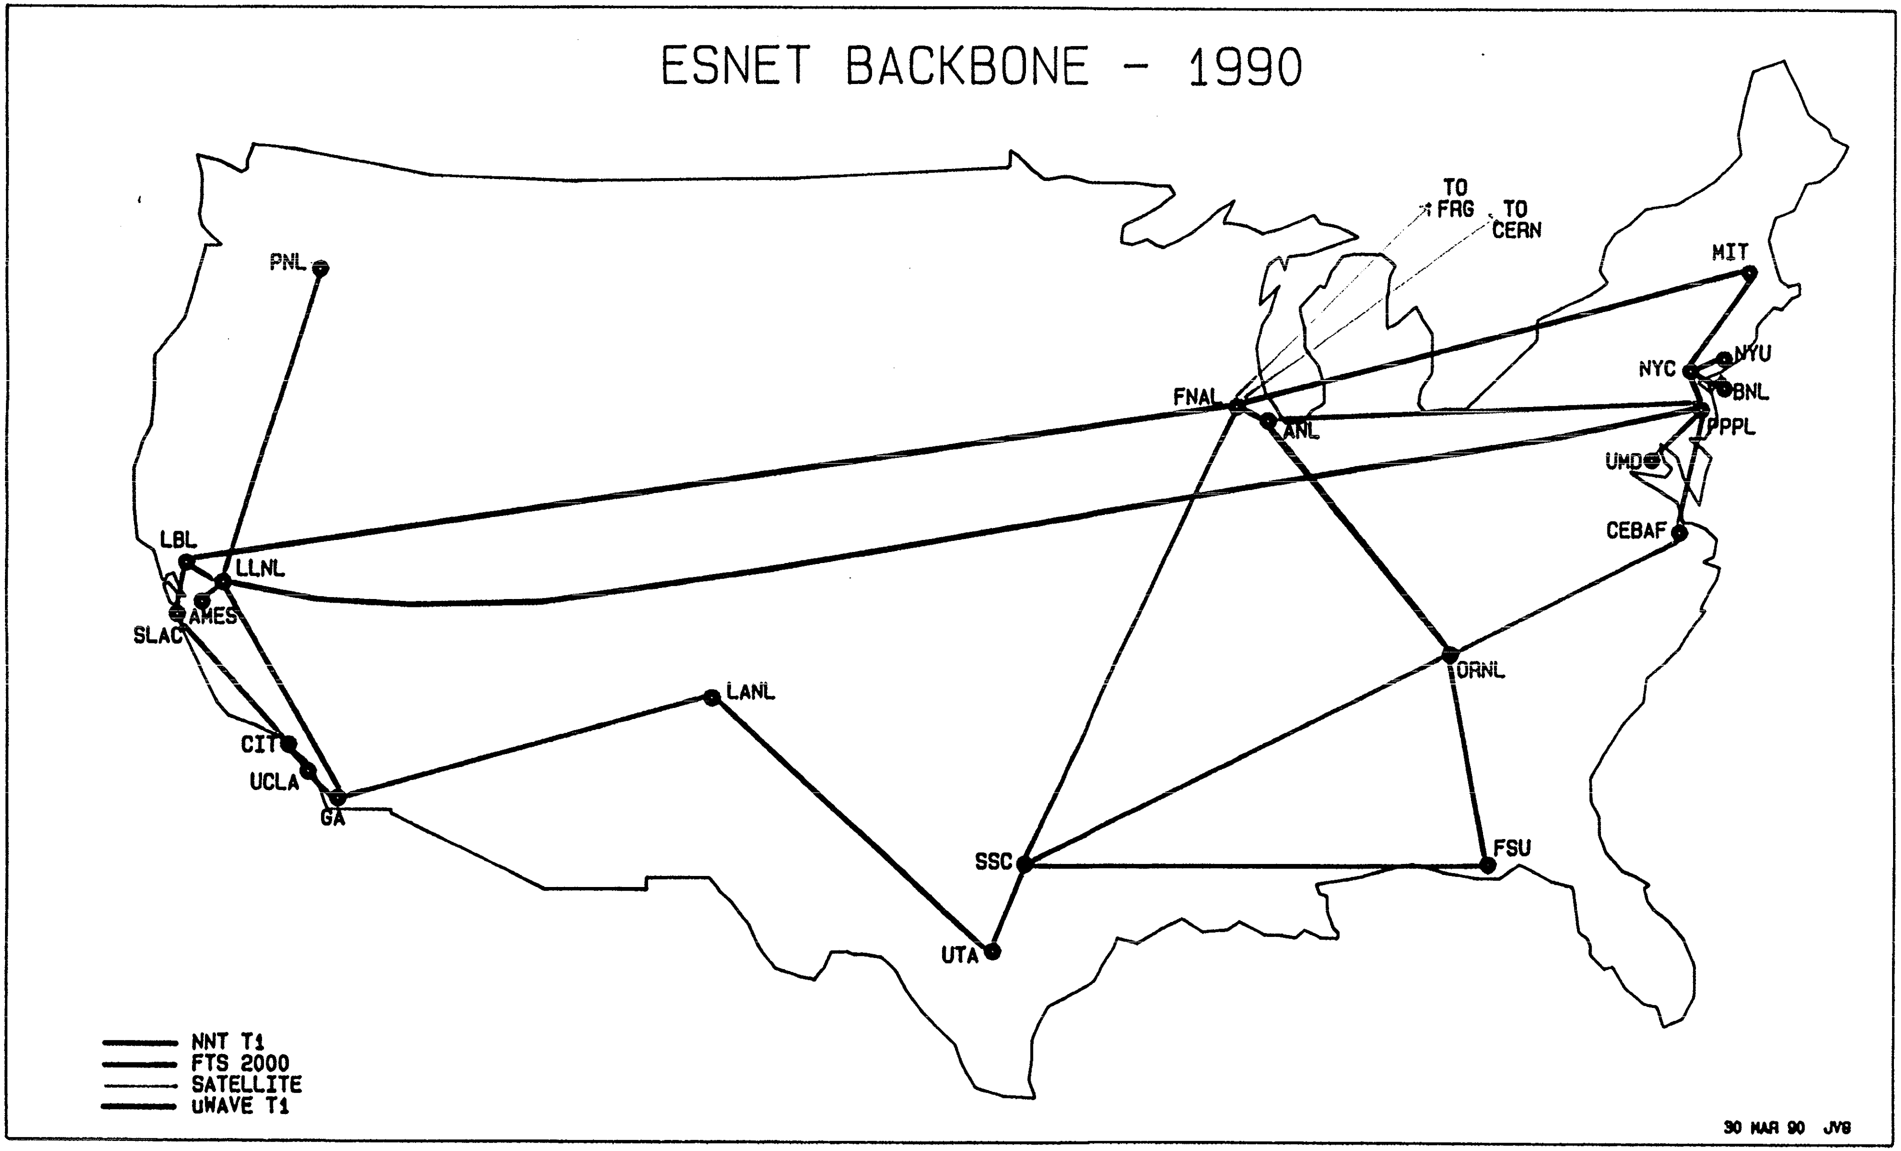

ESNET BACKBONE - 1990

ESnet

APRIL '90 STATUS

PAST ACTIVITIES:

DEBUG OF FTS-2000 LINES	FEB - MAR
MEETINGS WITH SITES AND REGIONALS	FEB - APR
PEERING WITH NSF/NSN/MB	MAR
PEERING WITH CONNECTED REGIONALS SURANET / BARRNET / THENET / SESQUINET / CERFNET / LOS NETTOS	APR
DECNET TRACEROUTE THROUGH CISCO's	MAR

ESnet

APRIL '90 STATUS

PLANED ACTIVITIES:

ADDITIONAL SITES
DOE / SNL-A&L / SAIC / AMES-IOWA

PEERING WITH ADDITIONAL REGIONALS
CICNET / NEARNET / NWNET / WESTNET /
MIDNET / JVNCNET

MOVE ITER / FRG LINE TO PPPL

UPGRADE CSC-1'S TO CSC-3'S

ESnet

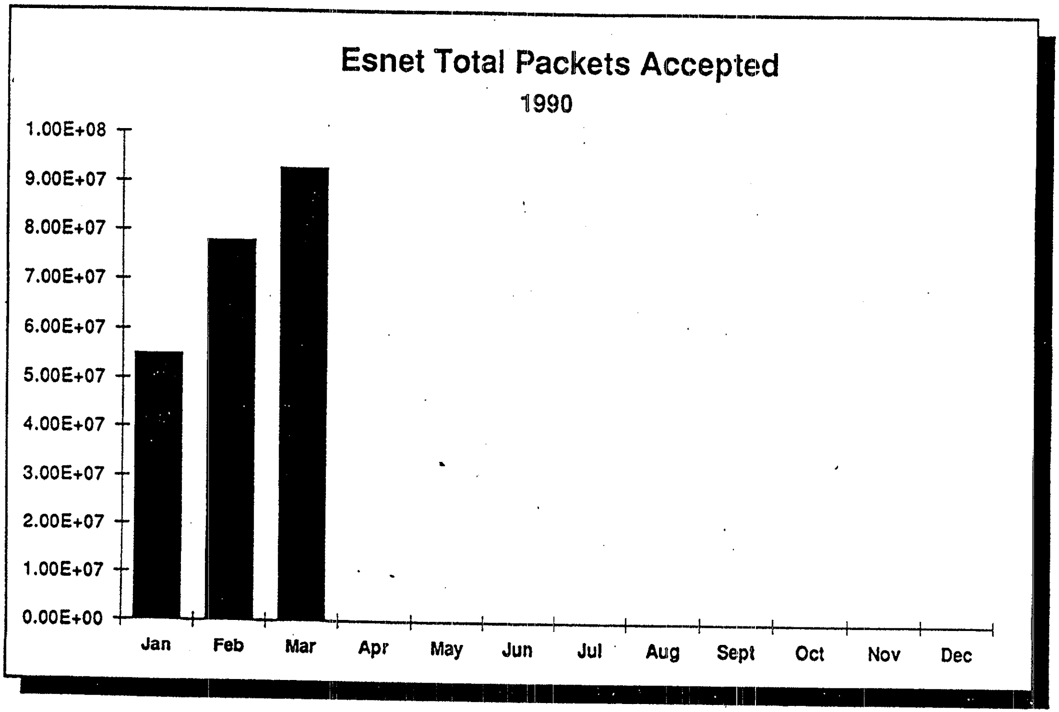
Esnet Total Packets Accepted
1990

ESnet

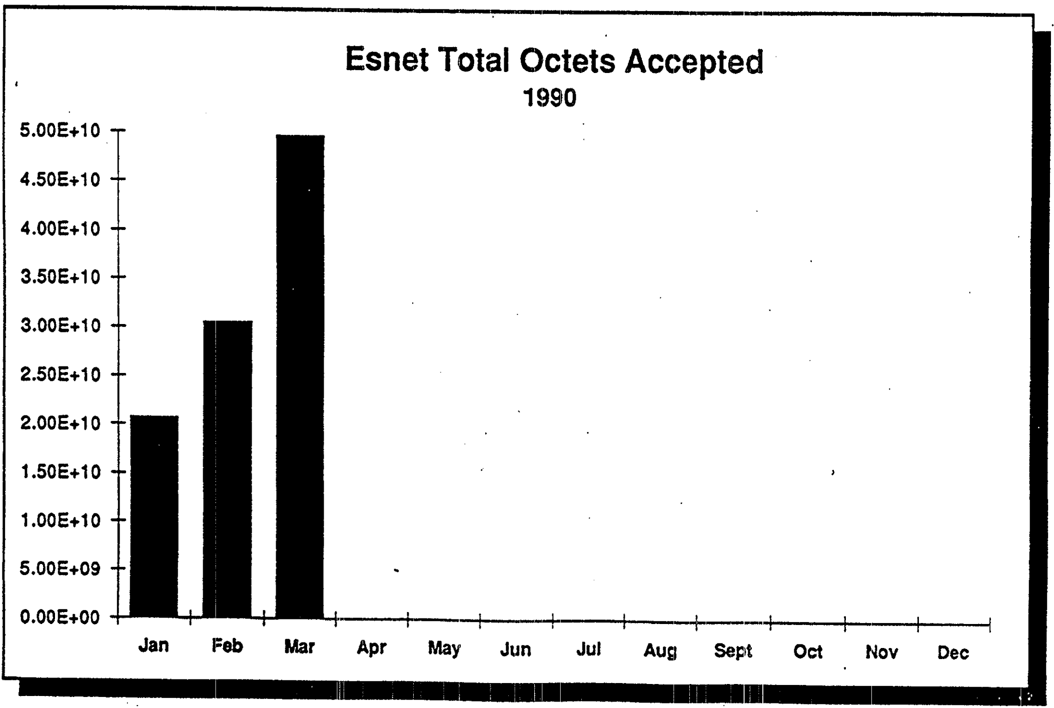
Esnet Total Octets Accepted
1990

4.3 NASA Sciences Internet Report

Presentation by Milo Medin/NASA

NASA Science Internet Report

Milo S. Medin

Sterling Software Corporation
NASA Science Internet Project Office
NASA Ames Research Center

Routing

- OSPF based, w/RIP and EGP at edges
- No non-OSPF transit routers
- FIX-net routers originate default
- Not used as a transit for regionals

NSI / NASA Science Network

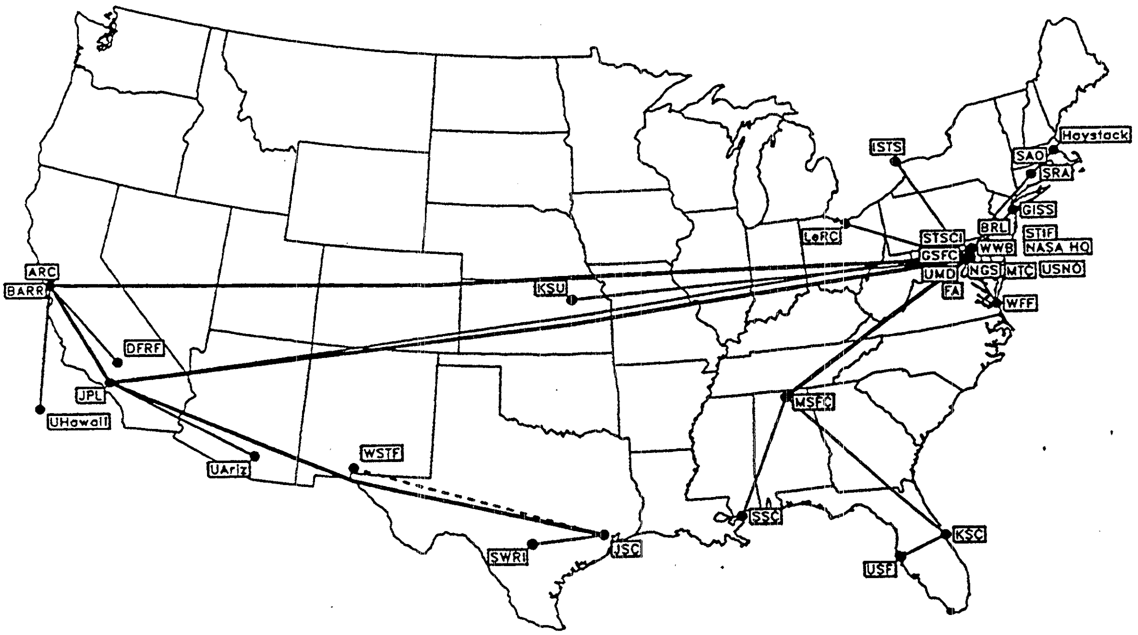

Prepared for: NASA Science Internet Project Office
by: Sterling Software, NAS2-11555
1 March 1990

NASA Science Network (NSN)

- 35 Routers (Proteon p4200)
- ~250 nets
- Multiprotocol (IP + DECNET)
- Regional and site connections
- International links
- Upcoming connections

Operational support

- NOC located at NASA Ames Research Center
- 24x7 staffing - 415-604-3655
- Out of band access to router consoles
- SNMP based tools
- New NOC under construction

OSPF deployment

- Switched on at 2355 4/13 (a Friday)
- No RIP compatibility attempted
- Cutover from east to west (NOC)
- Route tagging wonderful!
- Routing traffic overhead has dropped
- Rapid rerouting capability
- It works! And fairly well.

1 minute peak data

SSC-MSFC

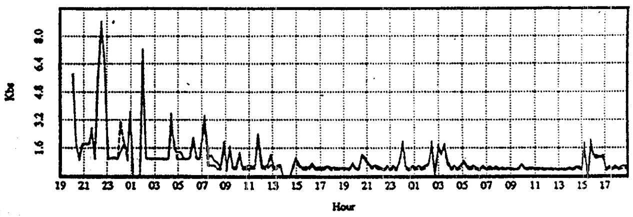

1 minute peak data

SSC-MSFC

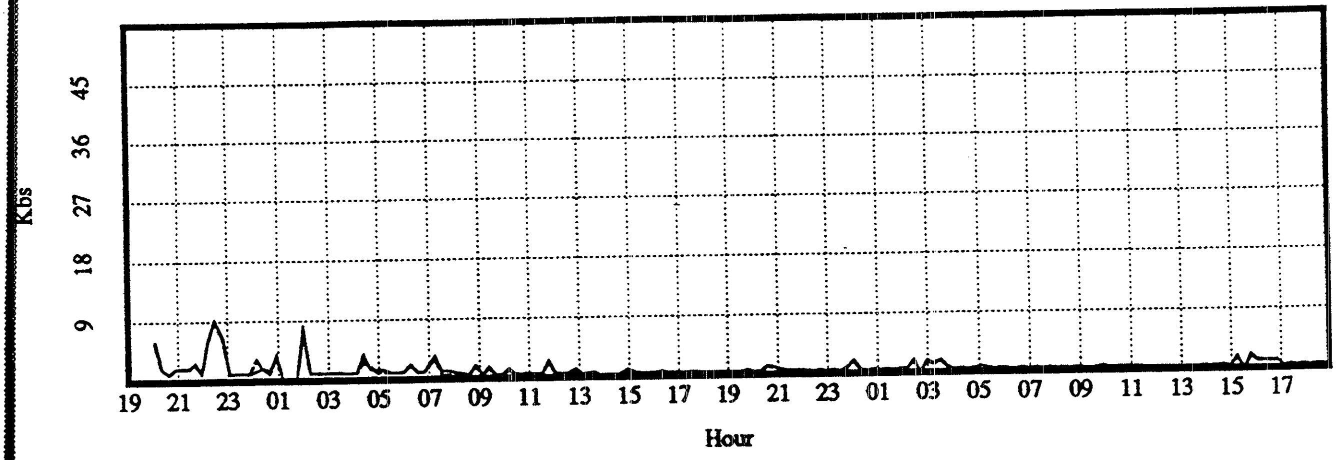

15 minute average data

SSC-MSFC

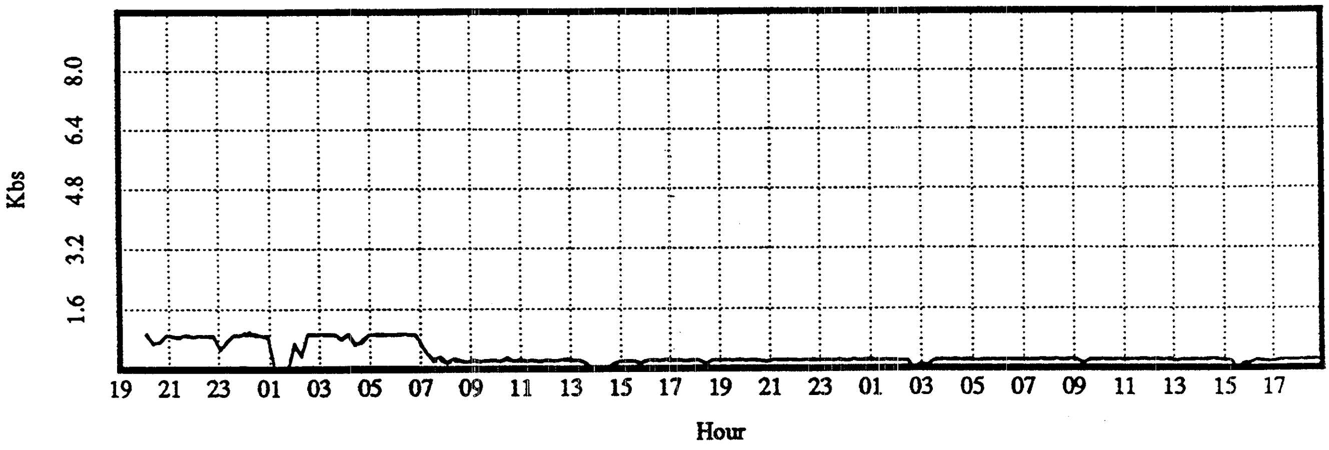

15 minute average data

SSC-MSFC

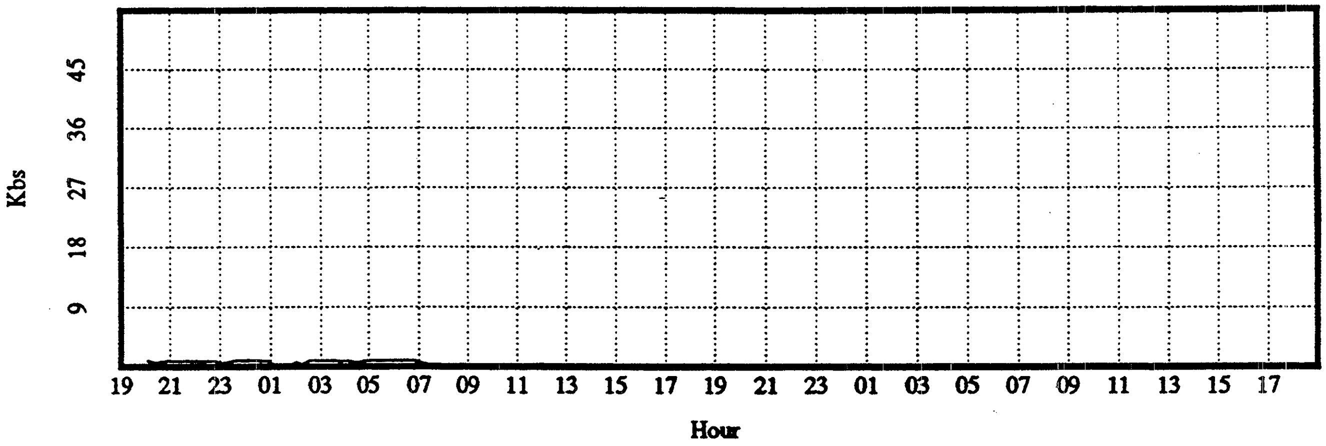

```
cincsac [1767]: ping nssdca.gsfc.nasa.gov
PING nssdca.gsfc.nasa.gov (128.183.10.4): 56 data bytes
64 bytes from 128.183.10.4: icmp_seq=1 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=2 time=99 ms
64 bytes from 128.183.10.4: icmp_seq=3 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=4 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=5 time=119 ms
64 bytes from 128.183.10.4: icmp_seq=6 time=120 ms
64 bytes from 128.183.10.4: icmp_seq=7 time=120 ms
64 bytes from 128.183.10.4: icmp_seq=8 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=9 time=120 ms
64 bytes from 128.183.10.4: icmp_seq=10 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=11 time=140 ms
64 bytes from 128.183.10.4: icmp_seq=12 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=13 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=14 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=15 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=16 time=100 ms   <-- Link shutdown
64 bytes from 128.183.10.4: icmp_seq=19 time=100 ms   <-- Path recovery
64 bytes from 128.183.10.4: icmp_seq=20 time=100 ms
64 bytes from 128.183.10.4: icmp_seq=21 time=120 ms
64 bytes from 128.183.10.4: icmp_seq=22 time=99 ms
64 bytes from 128.183.10.4: icmp_seq=23 time=100 ms
.
.
.
```

4.4 NSF Report

Presentation by Hans-Werner Braun/MERIT

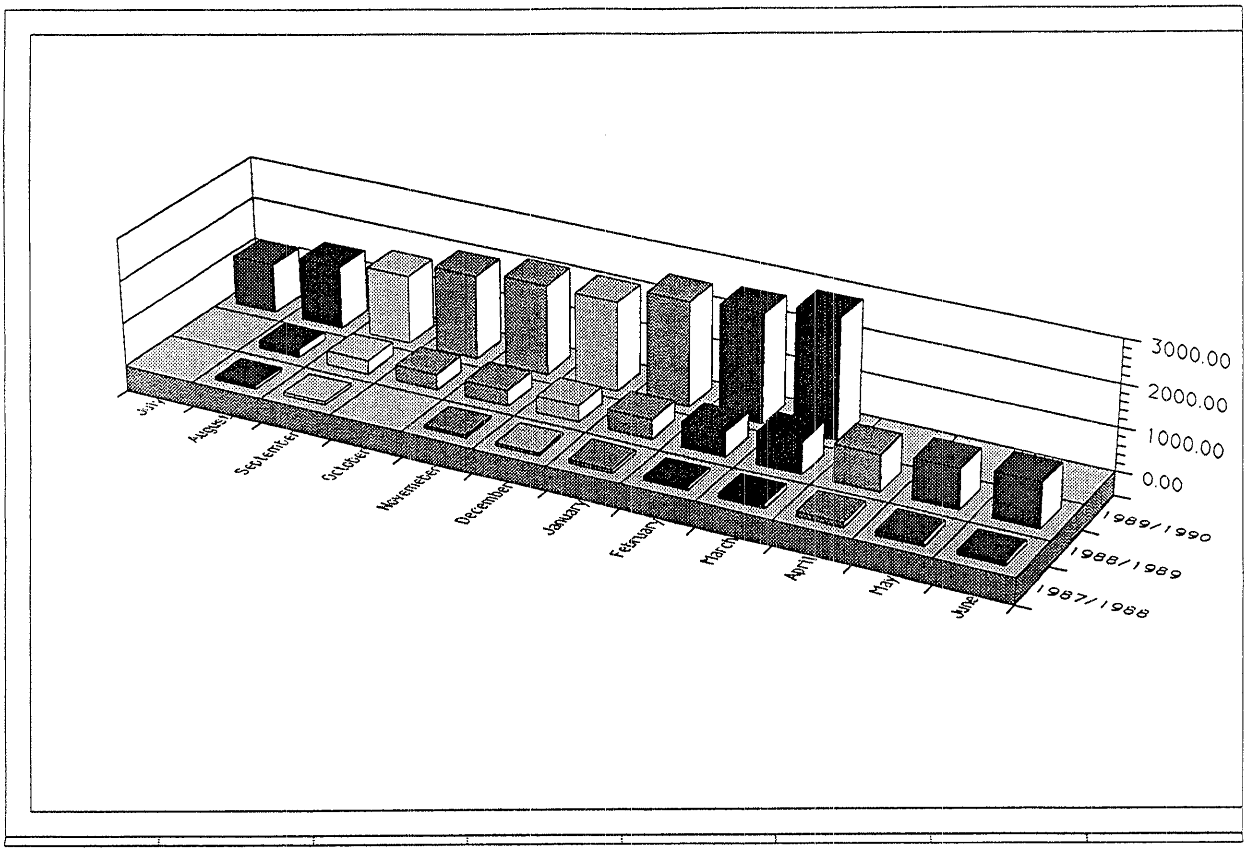
July
August
September
October
Novemeber
December
January
February
March
April
May
June
3000.00
2000.00
1000.00
0.00
1989/1990
1988/1989
1987/1988

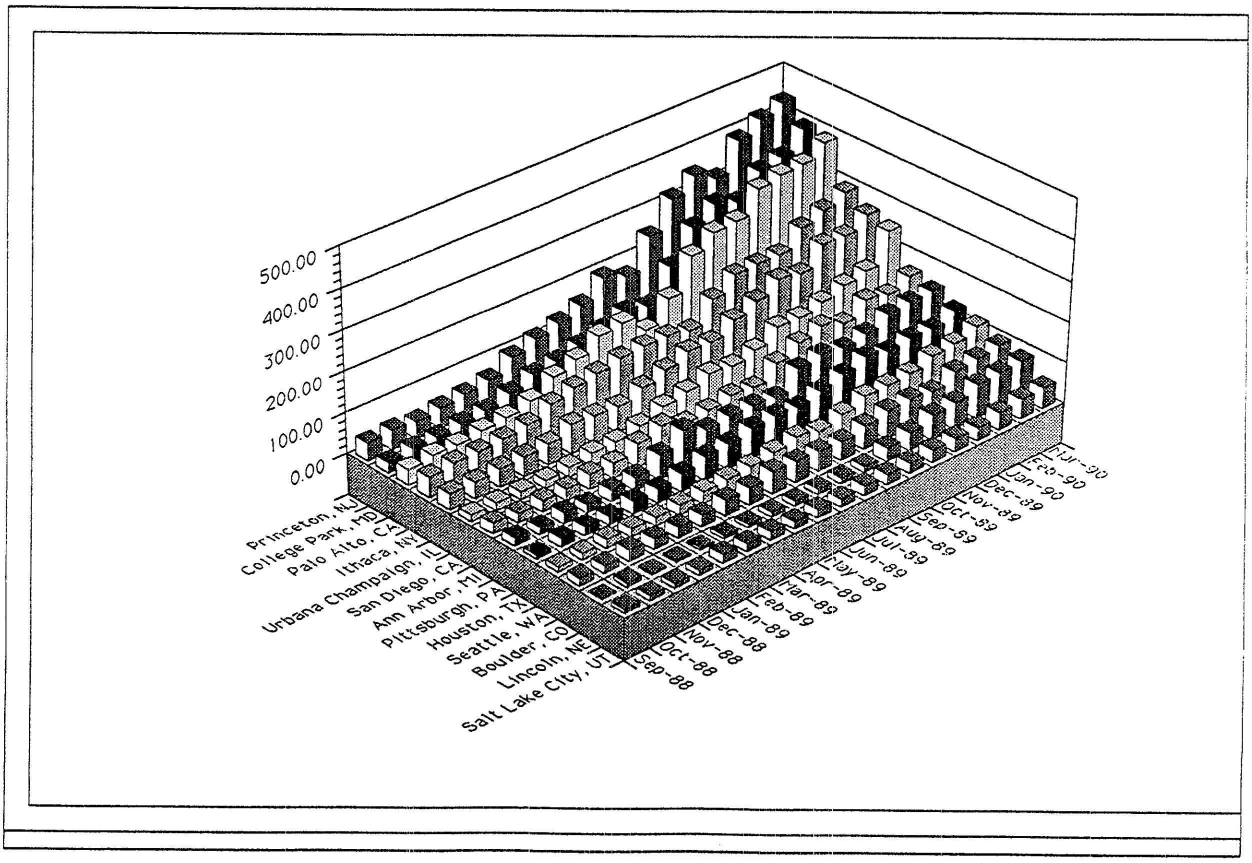
500.00
400.00
300.00
200.00
100.00
0.00
Princeton, NJ
College Park, MD
Palo Alto, CA
Ithaca, NY
Urbana Champaign, IL
San Diego, CA
Ann Arbor, MI
Pittsburgh, PA
Houston, TX
Seattle, WA
Boulder, CO
Lincoln, NE
Salt Lake City, UT
Sep-88
Oct-88
Nov-88
Dec-88
Jan-89
Feb-89
Mar-89
Apr-89
May-89
Jun-89
Jul-89
Aug-89
Sep-89
Oct-89
Nov-89
Dec-89
Jan-90
Feb-90
Mar-90

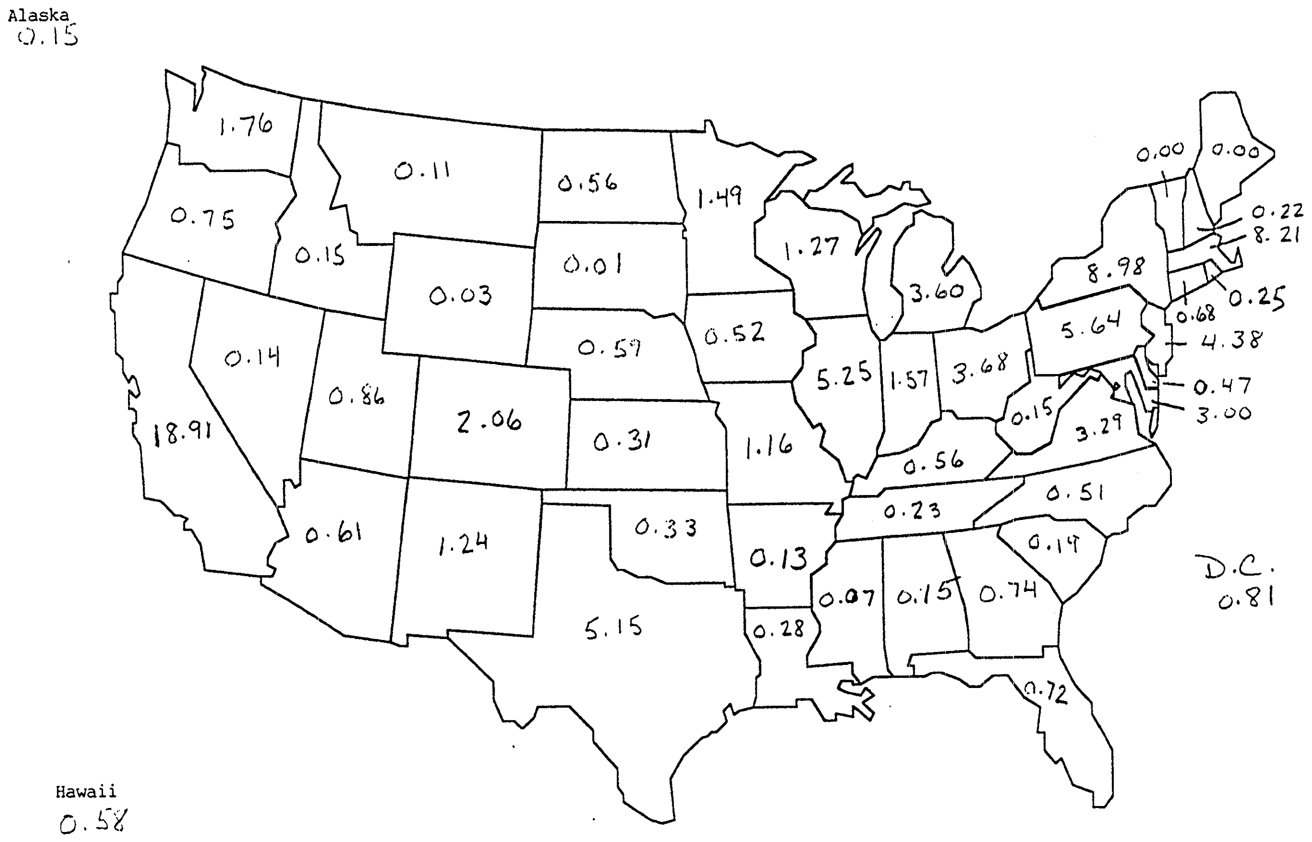

Percentage of packets in, November 1989

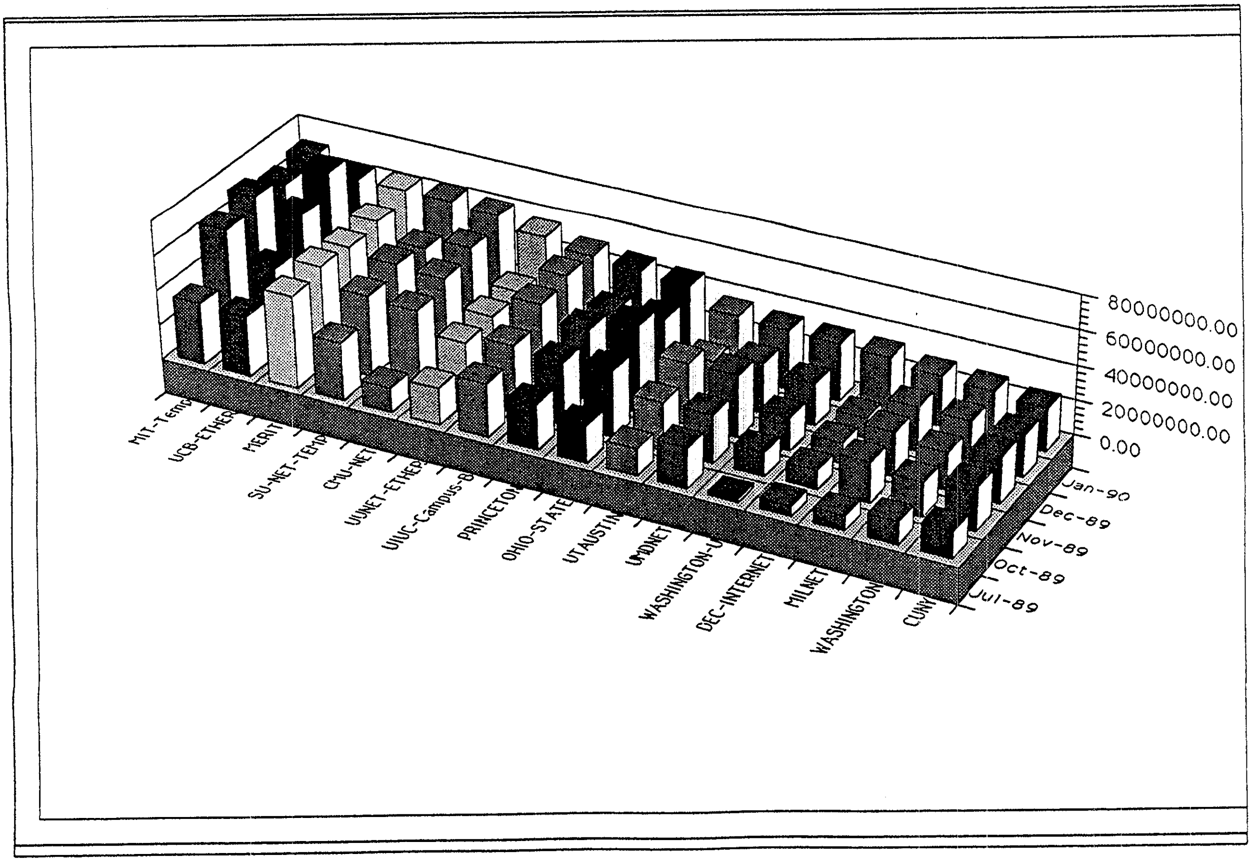

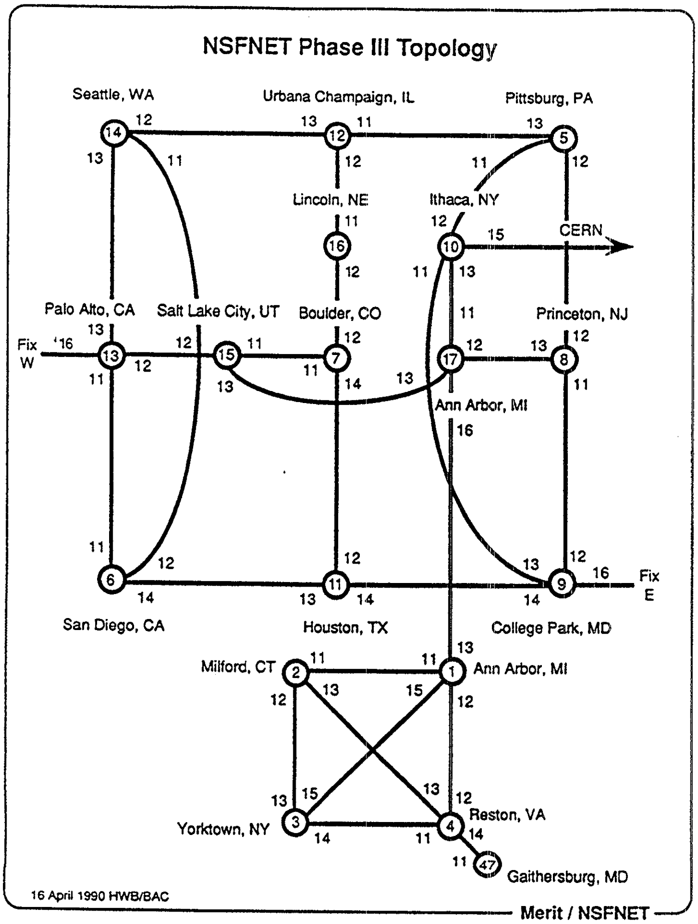

NSFNET Phase III Topology

16 April 1990 HWB/BAC

Merit / NSFNET

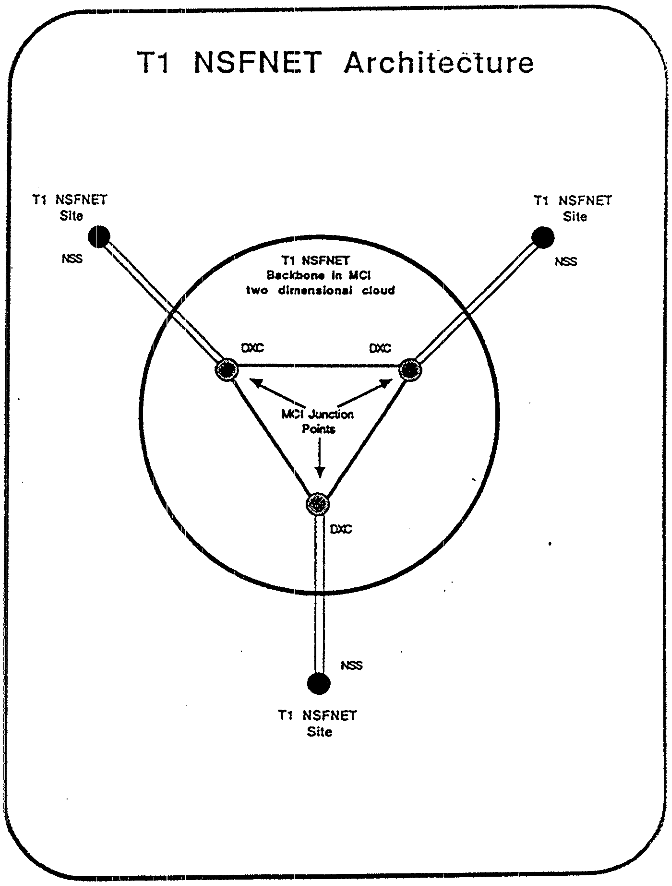

T1 NSFNET Architecture

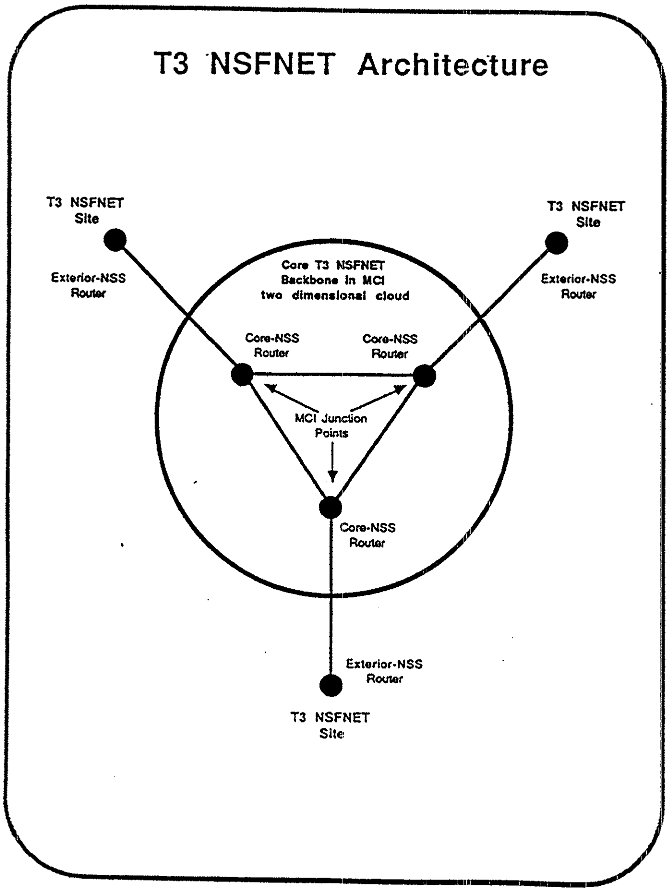

T3 NSFNET Architecture

Proposed NSFNET T3 Model

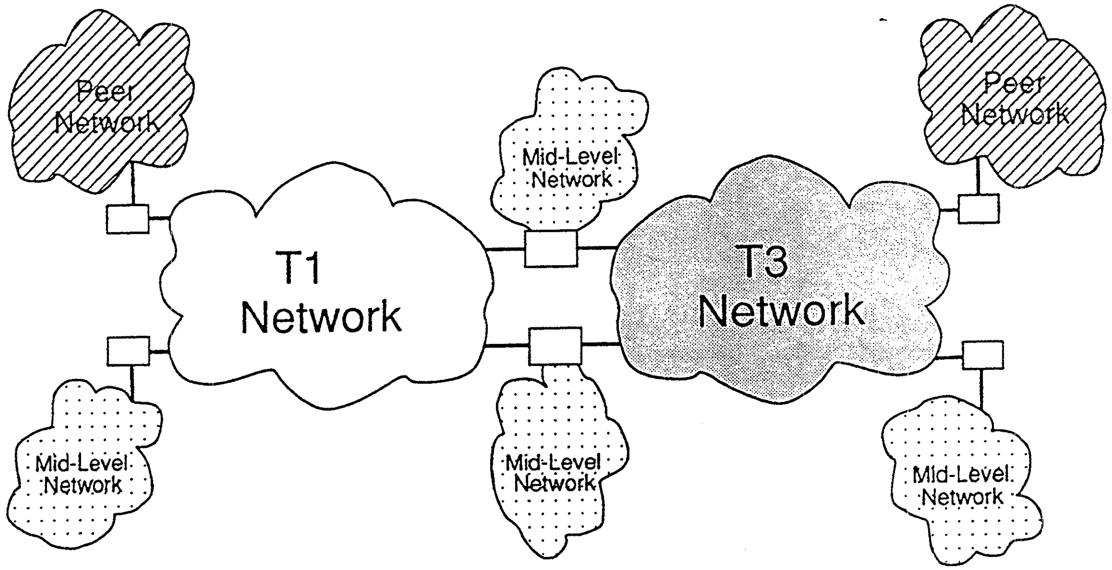

NSFNET

Early OSI Support in the NSFNET Backbone Network

- * InterOp '89 Demo
- * CLNP Support
- * Experimental Prototype
- * Early Experience on National Infrastructure

Qualifications of NOC Operators

1. BS Comp Science 6 mo exp user interfaces
2. 75% CS degree 8 years programming exp.
3. 75% EE degree, 2 AD 1 year help desk exp.
4. 75% CS degree 2 years Op exp. at chem plant in Brazil
5. 75% CS degree. 2.5 years Lab monitor/help desk
6. 75% CS degree 1 year programming
7. HS 7 years operation exp.
8. Assoc degree. CS
9. BFA 2 years lab tech
10. BS Anthro 2 years o
11. 2 years undergrad
12. HS 1 year operations/help desk

Posted: 6:14pm EDT, Wed Apr 25/90
Subject: NetOp "Extra" Activities
To: Hans-Werner Braun
From: Dale S. Johnson
Network information: none

HWB---

You had requested a list of some of the "extra" things NSFNet NetOpTechs do for the network. Here are a few:

- NSFNet Configuration
- Routing Problem debugging
- Tool Development
  - Configuration
  - Script Development (special temporary needs)
  - Simple tool development
  - Tool evaluation
- Database maintenance
- Unix system administration
- X-tool configuration

- NSS Assembly & Configuration
- Unix workstation assembly and configuration
- Cable tracing/debugging
- ROM replacement

- Mainframe backups (MVS, VM)
- Unix backups
- Backup procedure scripting

- Outage report preparation
- Statistical Reports on NOC Operation
- Presentation slide preparation

- Videotaping lectures, tape library maintenance

```
NSFNET performance report:  Sat Apr 28 1990

        Input/Output Traffic ( sorted by 15 minute peak samples )

                              INPUT
                              -----

                  Packets / sec.                             %        P/A
Ext-If   Min     Aver     Peak    Peak 15 min. period    of total   ratio   +    -
-----------------------------------------------------------------------------------

10-E0:  32.41    88.03   483.46   18:30:50 - 18:45:53     9.37     5.49    2    0
08-E0:  71.68   156.02   291.77   03:00:34 - 03:15:36    16.61     1.87    4    0
09-E0:  72.71   140.64   216.48   02:15:41 - 02:30:42    14.97     1.54    3    1
13-E0:  34.30   104.30   205.71   00:01:42 - 00:16:20    11.10     1.97    1    0
11-E0:  17.16    53.56   153.46   04:45:59 - 05:01:01     5.70     2.86    1    0
12-E0:  20.97    73.59   138.82   02:16:09 - 02:31:11     7.83     1.89    1    0
06-E0:  17.27    58.81   116.01   00:00:16 - 00:15:15     6.26     1.97    2    0
17-E0:   4.45    50.66    90.85   19:02:11 - 19:16:55     5.39     1.79    1    2
05-E0:  20.44    49.43    86.02   20:15:17 - 20:30:07     5.26     1.74    4    0
14-E0:  19.68    39.80    72.62   10:01:28 - 10:16:26     4.23     1.82    4    1
FIX-W:  20.34    33.45    58.77   01:01:20 - 01:16:17     3.56     1.76    4    0
16-E0:   8.93    30.40    54.63   21:46:52 - 22:01:59     3.23     1.80    1    0
07-E0:  12.35    28.67    51.47   00:15:27 - 00:30:27     3.05     1.80    2    0
15-E0:   2.81    12.58    30.85   00:16:39 - 00:31:40     1.34     2.45    4    0
FIX-E:   9.65    15.27    24.91   19:00:44 - 19:15:46     1.63     1.63    4    0
EASIN:   0.49     1.92    12.64   18:45:53 - 19:15:56     0.20     6.58    7    0
R-NET:   1.07     2.18     2.60   23:31:59 - 23:46:51     0.23     1.19    0    3
CNUSC:   0.09     0.21     1.49   14:45:50 - 15:00:49     0.02     7.02    2    0
-----------------------------------------------------------------------------------

                              OUTPUT
                              ------

                  Packets / sec.                             %        P/A
Ext-If   Min     Aver     Peak    Peak 15 min. period    of total   ratio   +    -
-----------------------------------------------------------------------------------

10-E0:  34.39    90.13   470.64   18:30:48 - 18:45:51     9.66     5.22    1    0
08-E0:  75.21   157.90   274.28   05:00:33 - 05:15:33    16.94     1.74    4    0
09-E0:  57.83   133.27   247.02   04:45:39 - 05:00:39    14.29     1.85    3    0
13-E0:  39.38   104.29   181.10   01:01:19 - 01:16:15    11.06     1.74    0    0
12-E0:  24.66    78.37   135.54   01:31:06 - 01:46:10     8.40     1.73    0    0
06-E0:  21.37    60.45   126.80   00:00:14 - 00:15:14     6.49     2.10    2    0
17-E0:  14.12    51.47   104.96   19:02:10 - 19:16:54     5.51     2.04    3    1
05-E0:  19.76    49.35    99.67   02:45:05 - 03:00:05     5.30     2.02    3    0
11-E0:  24.44    48.34    79.35   00:16:02 - 00:30:59     5.18     1.64    3    0
14-E0:  21.32    40.27    68.18   10:01:27 - 10:16:25     4.32     1.69    3    2
16-E0:   9.59    28.90    53.35   21:46:50 - 22:01:58     3.10     1.85    1    0
15-E0:  11.20    23.95    52.42   19:46:50 - 20:01:42     2.57     2.19    6    0
FIX-W:   8.94    22.77    50.06   22:31:34 - 22:46:21     2.42     2.20    4    0
07-E0:  12.93    29.44    45.82   23:15:28 - 23:30:28     3.16     1.56    0    0
FIX-E:   1.13    11.05    24.85   18:45:42 - 19:00:43     1.19     2.25    6    0
EASIN:   0.49     1.66     8.74   19:00:51 - 19:15:55     0.18     5.27    7    0
R-NET:   0.95     2.14     2.71   23:31:59 - 23:46:51     0.23     1.27    2    2
CNUSC:   0.13     0.21     1.21   14:45:48 - 15:00:48     0.02     5.74    2    0

Total Input: 8.111626e+07 packets
Total Output: 8.052365e+07 packets
```

```
        Link Utilization ( sorted by 15 minute peak utilization )

 Node         Utilization %                                     P/A
src -> dst  min     aver     max       15 min. peak         ratio    +    -
------------------------------------------------------------------------------
09 -> 11    4.05    15.20    29.51    00:15:51 - 00:30:50   1.94     1    1
11 -> 09    5.90    14.58    26.91    01:16:04 - 01:31:05   1.85     1    3
05 -> 08    3.29    10.40    25.30    02:45:12 - 03:00:12   2.43     4    0
08 -> 09    3.58    9.90     22.97    20:15:52 - 20:30:40   2.32     5    0
13 -> 15    2.34    8.01     19.96    04:16:42 - 04:31:29   2.49     5    0
06 -> 14    1.02    4.88     18.22    16:30:23 - 16:45:33   3.73     3    0
13 -> 06    1.04    5.06     17.80    14:01:32 - 14:16:32   3.52     3    0
17 -> 08    1.38    4.82     17.42    20:17:21 - 20:32:01   3.61     6    0
15 -> 17    2.47    6.41     17.28    04:16:59 - 04:31:46   2.70     7    0
12 -> 05    0.68    8.33     17.18    03:46:14 - 04:01:16   2.06     1    1
15 -> 13    2.18    5.77     16.84    14:01:49 - 14:16:49   2.92     4    0
17 -> 15    2.12    5.57     16.55    14:02:02 - 14:17:02   2.97     3    0
11 -> 06    1.78    6.04     15.79    00:16:09 - 00:31:09   2.61     4    0
12 -> 14    1.61    4.67     15.16    01:31:13 - 01:46:16   3.25     3    0
08 -> 05    2.47    8.40     14.73    21:15:43 - 21:30:37   1.75     1    0
17 -> 10    0.84    2.79     14.65    04:17:16 - 04:32:04   5.25     4    0
14 -> 12    0.77    5.04     14.39    02:16:35 - 02:31:39   2.85     5    0
10 -> 09    0.40    2.11     13.88    02:15:56 - 02:30:57   6.58     3    0
14 -> 13    1.41    3.57     13.88    01:31:36 - 01:46:36   3.89     3    0
05 -> 12    0.78    6.89     13.65    02:15:12 - 02:30:13   1.98     3    1
15 -> 07    1.44    3.82     13.42    19:46:55 - 20:01:47   3.52     5    0
09 -> 08    2.42    6.85     12.85    19:00:50 - 19:15:53   1.88     3    0
06 -> 13    1.71    6.05     12.50    21:00:27 - 21:15:23   2.07     1    0
07 -> 11    1.37    2.91     11.79    19:45:30 - 20:00:33   4.05     5    0
07 -> 16    0.57    3.29     11.63    14:30:30 - 14:45:31   3.54     2    0
08 -> 17    0.44    5.80     11.62    00:00:41 - 00:15:40   2.00     3    1
16 -> 12    0.46    2.58     11.04    14:32:01 - 14:46:50   4.28     3    0
05 -> 10    1.14    5.25     10.54    15:45:22 - 16:00:12   2.01     2    1
13 -> 14    1.14    3.65     10.19    02:16:25 - 02:31:28   2.79     5    0
10 -> 05    1.80    4.97     9.65     20:01:00 - 20:16:17   1.94     4    0
12 -> 16    0.61    2.32     9.06     00:16:17 - 00:31:18   3.90     6    0
11 -> 07    0.41    2.63     8.56     14:31:04 - 14:46:05   3.25     1    0
16 -> 07    0.48    2.28     8.55     00:16:50 - 00:32:00   3.75     5    0
07 -> 15    0.51    1.67     6.22     17:15:31 - 17:30:33   3.72     2    0
10 -> 17    0.74    2.05     4.48     05:00:57 - 05:15:57   2.18     2    0
14 -> 06    0.27    1.82     4.46     20:16:58 - 20:31:41   2.44     3    0
06 -> 11    0.33    1.40     3.87     01:00:23 - 01:15:21   2.76     4    0
09 -> 10    0.34    1.67     3.87     02:45:48 - 03:00:49   2.32     5    0
```

Note:

1) All indicated times are in Universal Time (WET)

EST = WET - 4/5 hours (with Daylight Savings)

2) Input/Output traffic is a count of packets entering/leaving the NSFNET backbone system.

3) Link Utilization is defined as:

Utilization = bits transmitted / 1.344 Mb/sec (available bandwidth)

4) Ext-If: NSS interface to regional/peer networks.

5) min/average/max

The minimum resource utilization over all 15 minute periods.
The average resource utilization over the whole day.
The maximum resource utilization over all 15 minute periods.

6) P/A ratio: Peak (maximum) to Average ratio.

7) Columns labelled + and - :

This is an indicator of the degree of burstiness of the measured quantity, over the period of the given day.

+ : number of 15 minute samples greater than (average + 2*standard deviation)
- : number of 15 minute samples less than (average - 2*standard deviation)

For example, if + is 1, then there was probably only one large peak, and this peak could be as a result of a some event outside the "normal" pattern for that resource utilization (eg. large file transfer)

Node key:

05: Pittsburg
06: San Diego
07: Boulder
08: Princeton
09: College Park
10: Ithaca
11: Houston
12: Urbana-Champaign
13: Palo Alto
14: Seattle
15: Salt Lake City
16: Lincoln
17: Ann Arbor

Interface key:

E0: Directly attached ethernet interface
FIX-E and FIX-W: Federal Interagency eXchange
CNUSC: T1 link to CNUSC, France
EASIN: T1 link to CERN, Switzerland
R-NET: T1 connection to the Research Network

Chapter 5

IETF Protocol Presentations

5.1 ST-2: The Experimental Internet Stream Protocol

Presentation by Claudio Topolcic BBN

ST-2 includes some relatively minor improvements to ST, as described in IEN 119. ST is an internet protocol at the same layer as IP. ST differs from IP in that IP, as originally envisioned, did not require gateways or intermediate systems to maintain state information describing the streams of packets flowing through them. ST incorporates the concept of streams across the Internet. Every intervening ST entity maintains state information for each stream that passes through it. The stream state includes forwarding information, which efficiently supports multicast, as well as resource information, allowing networks or link bandwidths and queues to be managed for a specific stream. This pre-allocation of resources allows data packets to be forwarded with low delay and overhead, and with low probability of loss due to congestion. The characteristics of a stream, such as the number and location of the endpoints, and the bandwidth required, may be modified during the lifetime of the stream.

A stream is "multi-destination simplex" (MDS) since data travels across it in only one direction: from the origin to the targets. A stream can be viewed as a directed tree in which the origin is the root, all the branches are directed away from the root toward the targets, which are the leaves. Each stream is identified by a globally unique "NAME" that includes the origin's address plus a field to make the name unique. The name is specified in control operations, but is not used in ST data packets. ST data packets instead, contain a short hop-by-hop identifier (HID) which is used for efficient forwarding.

Host and gateway "ST agents" create and manage a stream by exchanging control messages. The control protocol follows a request-response model with retransmission after timeout. Control messages are used to: create streams, refuse creation of a stream, delete a stream in whole or in part, negotiate or change a stream's parameters, tear down parts of streams as a result of gateway or network failures, or transient routing inconsistencies, and reroute around network or component failures. ST stream setup control messages also carry a "next protocol identifier" which allows ST to demultiplex streams to a number of possible higher level protocols and a "port identifier" to further demultiplex to a specific instance of an application.

A stream is monitored by the involved ST agents. If they detect a failure, they can attempt a repair, which in general involves tearing down part of the stream and rebuilding it to bypass the failed component(s). If an ST agent determines it cannot effect the repair, it propagates the failure information back to its previous-hop agent. ST agents use a flow specification, called the "FlowSpec", to describe the required characteristics of a stream. Included are values for bandwidth, delay, and reliability parameters. The FlowSpec describes both the desired values as well as their minimal allowable values. The intervening ST agents thus have some freedom in allocating their resources. Intervening ST agents accumulate statistics that describe the characteristics of the chosen path and pass that information to the origin and the targets of the stream.

The Internet Stream Protocol (ST)

Claudio Topolcic

BBN Systems and Technologies Corporation

Engineer's Project

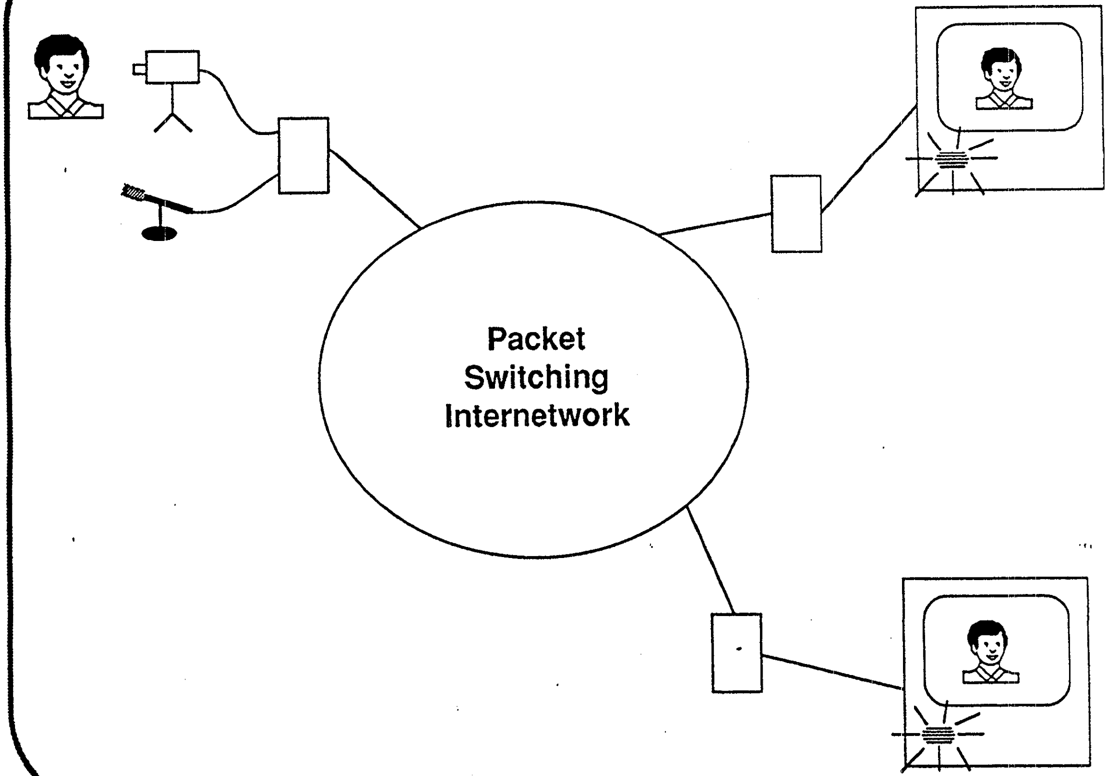

Packet Switching Internetwork

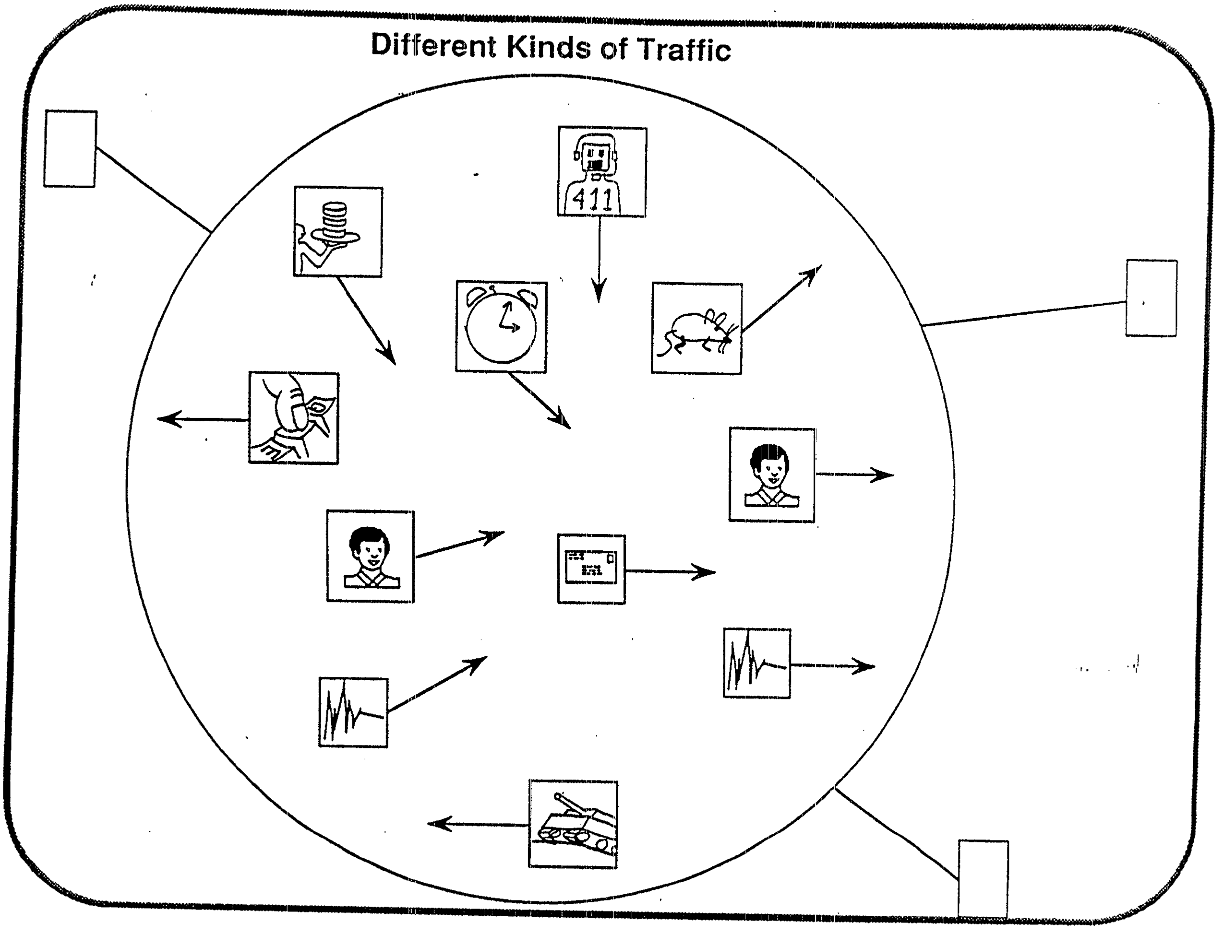

Different Kinds of Traffic

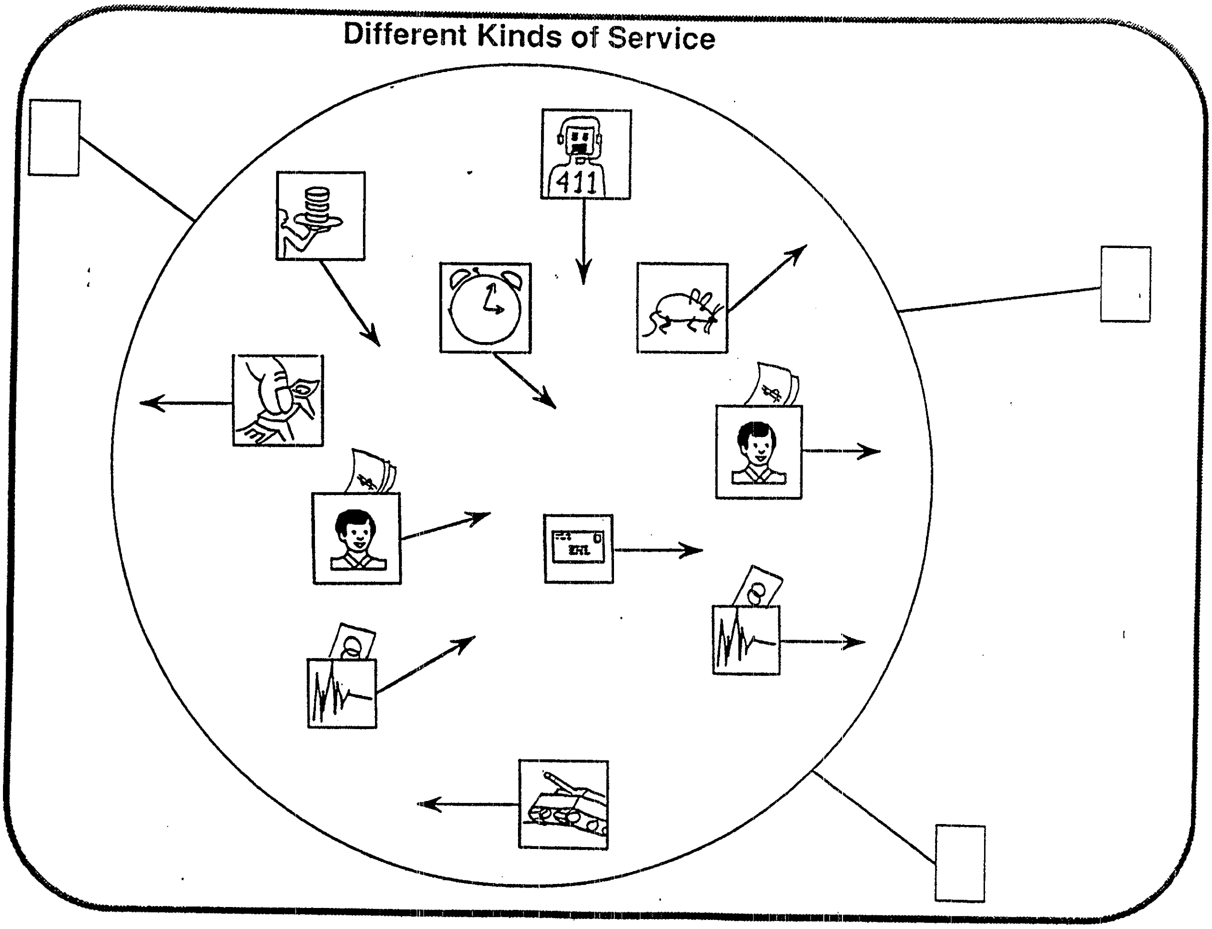

Different Kinds of Service

"Engineer's" Approach

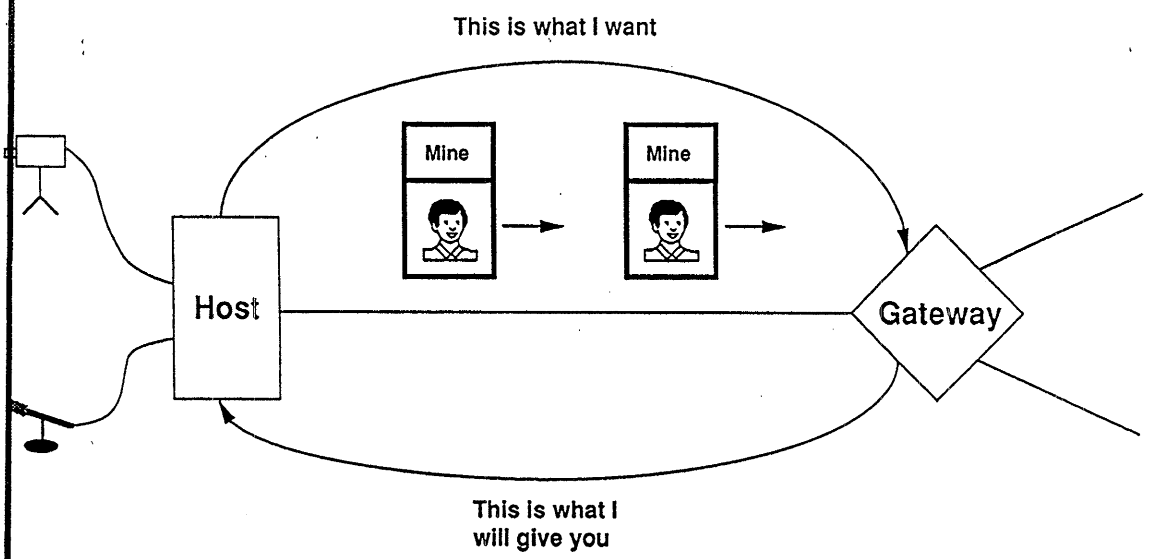
This is what I want

This is what I will give you

IP Model

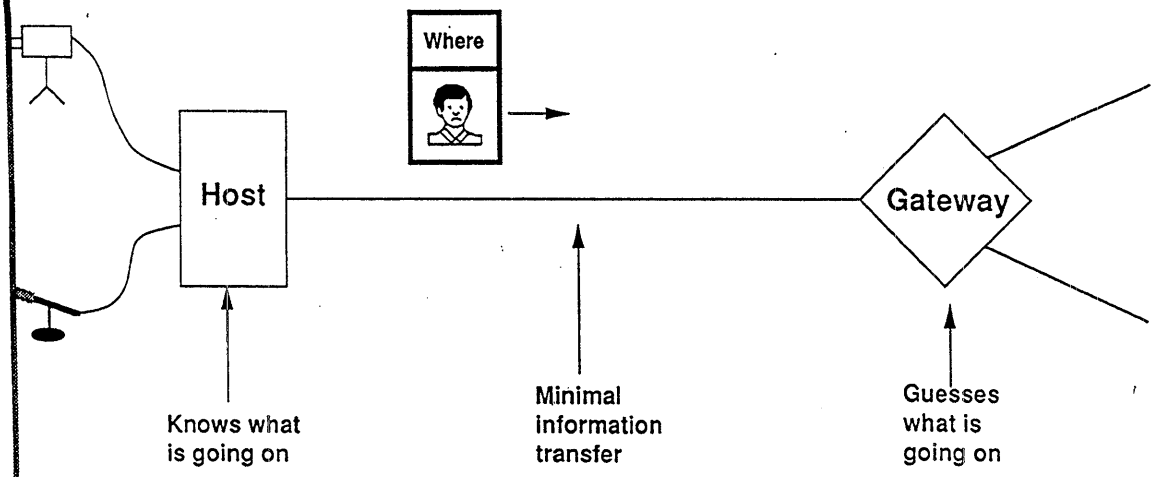
Knows what is going on

Minimal information transfer

Guesses what is going on

Engineering Optimization

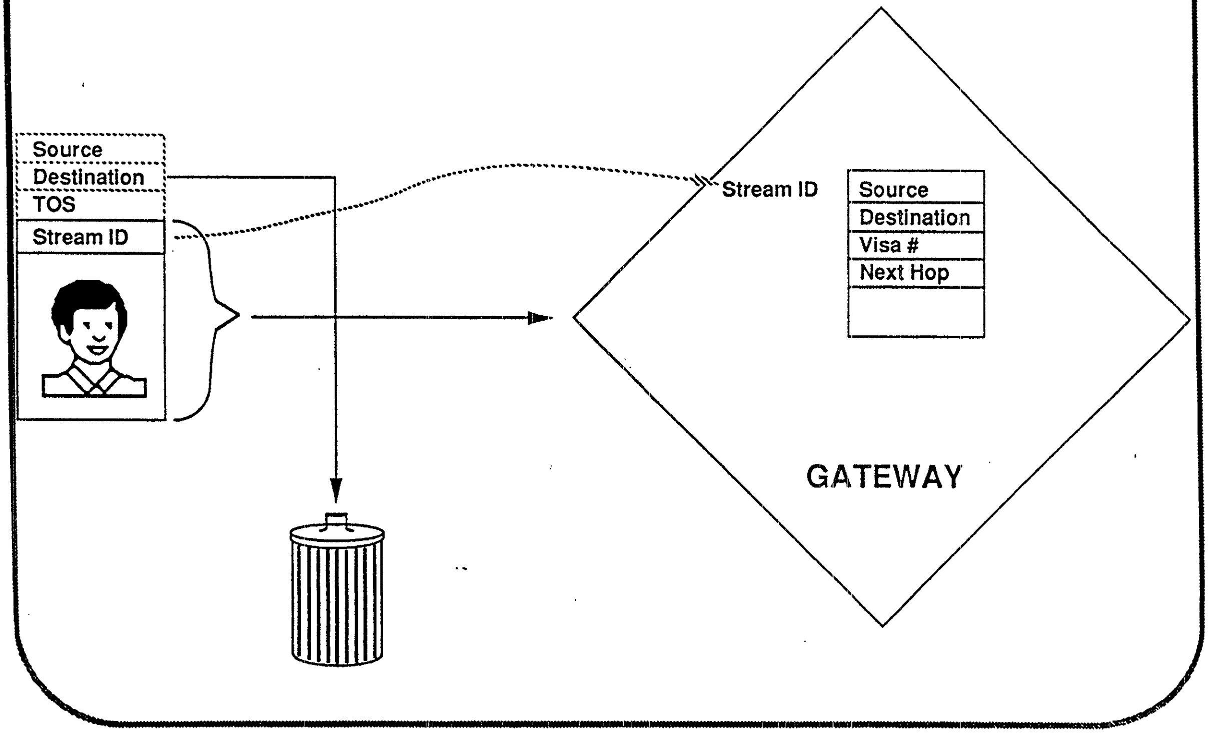

Source
Destination
TOS
Stream ID

Stream ID

Source
Destination
Visa #
Next Hop

GATEWAY

Engineering Tradeoffs

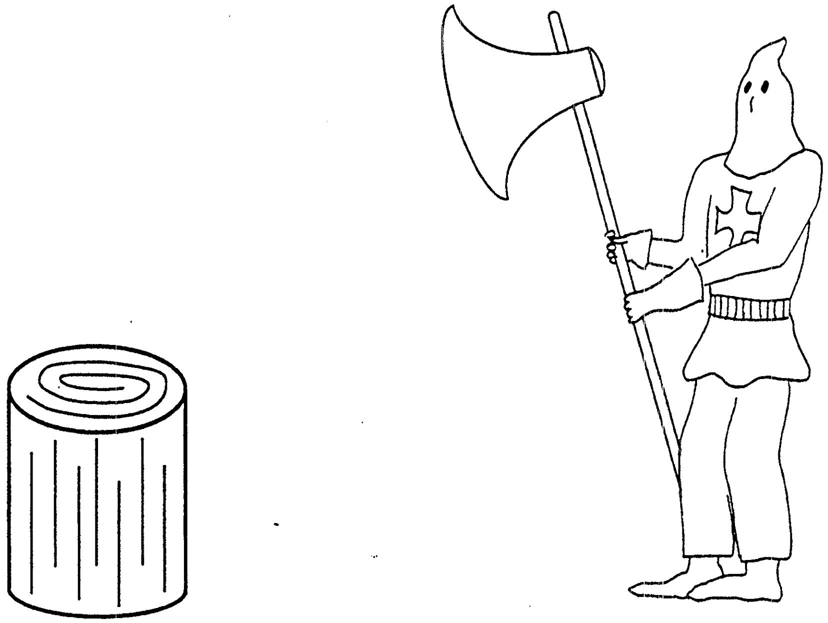

Some Gateway Functions

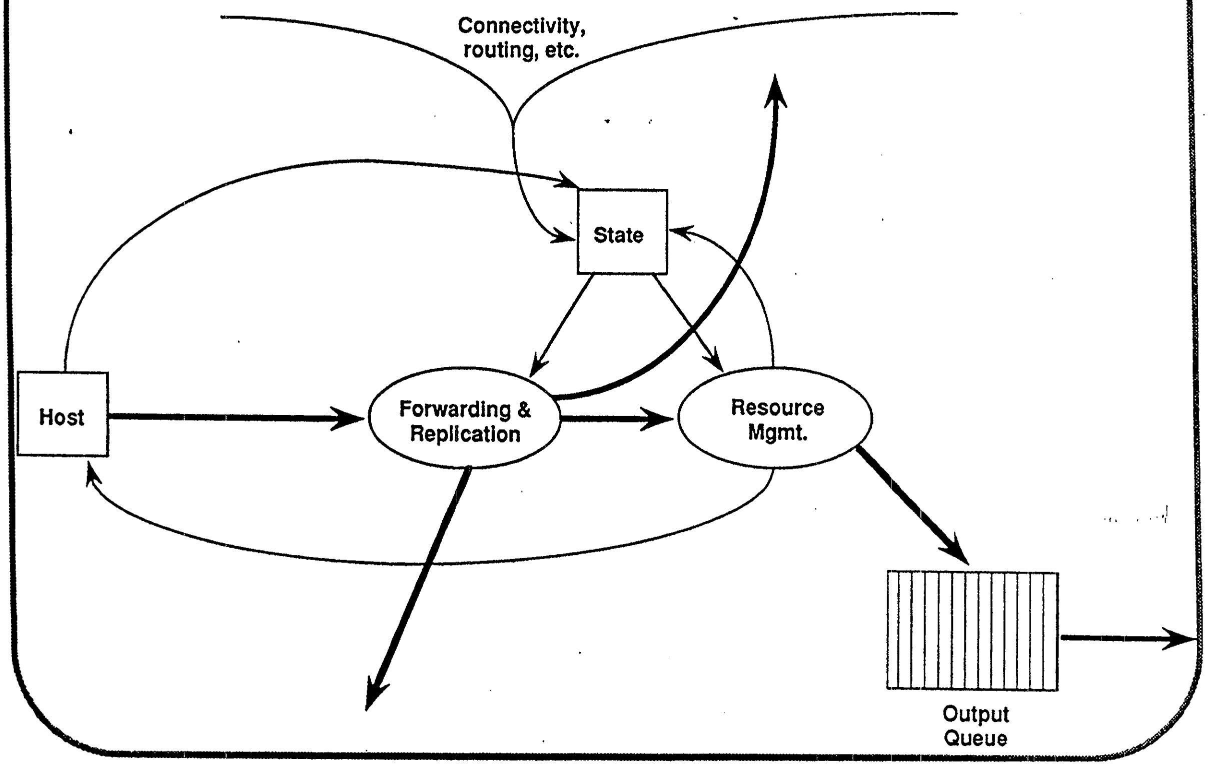

Connectivity, routing, etc.

State

Host

Forwarding & Replication

Resource Mgmt.

Output Queue

The "only interesting" part of the problem

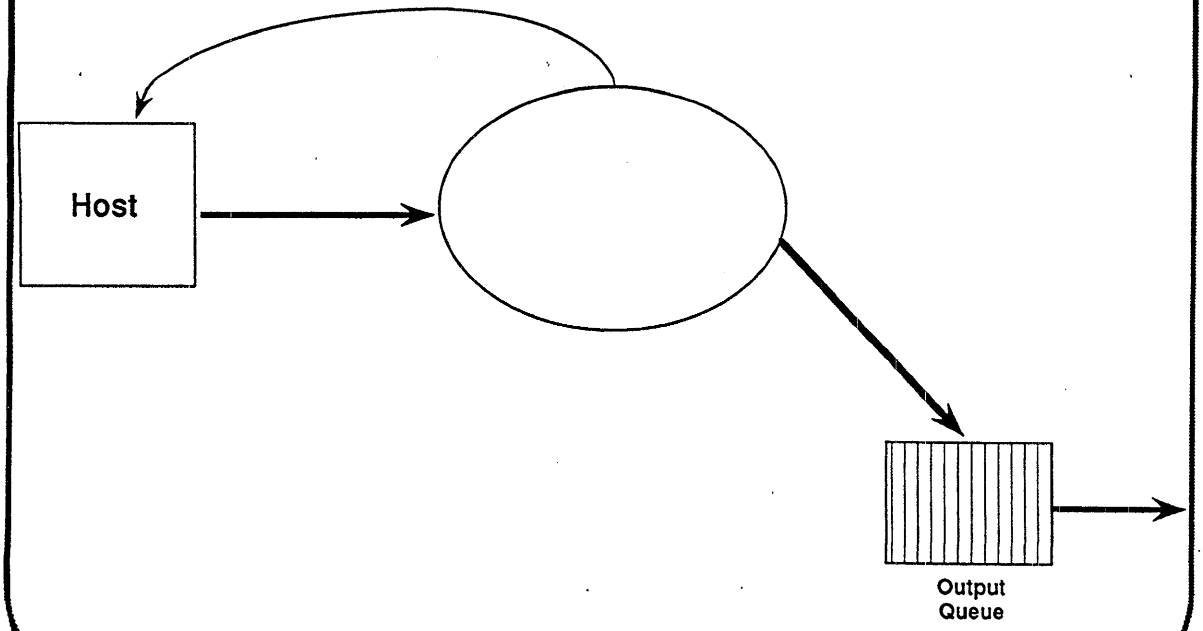

Host

Output Queue

Connection Oriented Protocol Functions

- Connection Management
- Connection Maintenance
- Routing
- Forwarding
- Multicasting
- Resource Management
  (The only interesting part of the protocol)
- Interoperation With Other Systems
- Security

The Internet Stream Protocol (ST)

- A Stream represents managed resources along a multicast path
- The Stream is the fundamental building block
- Structured as a directed tree
- Streams can be dynamically modified
- Allow minimum forward processing for data packets

Structure of a Stream

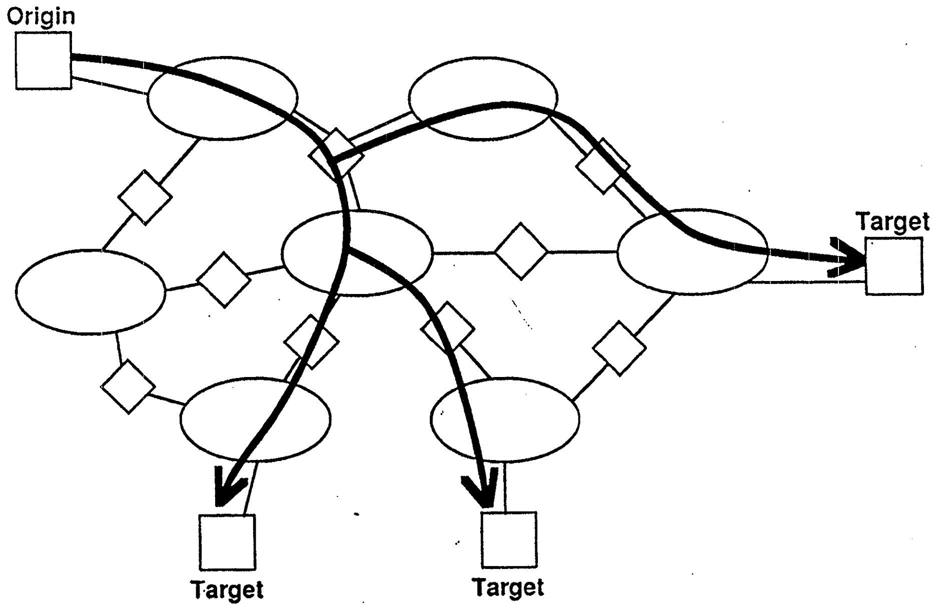

A Conference Built of Multiple Streams

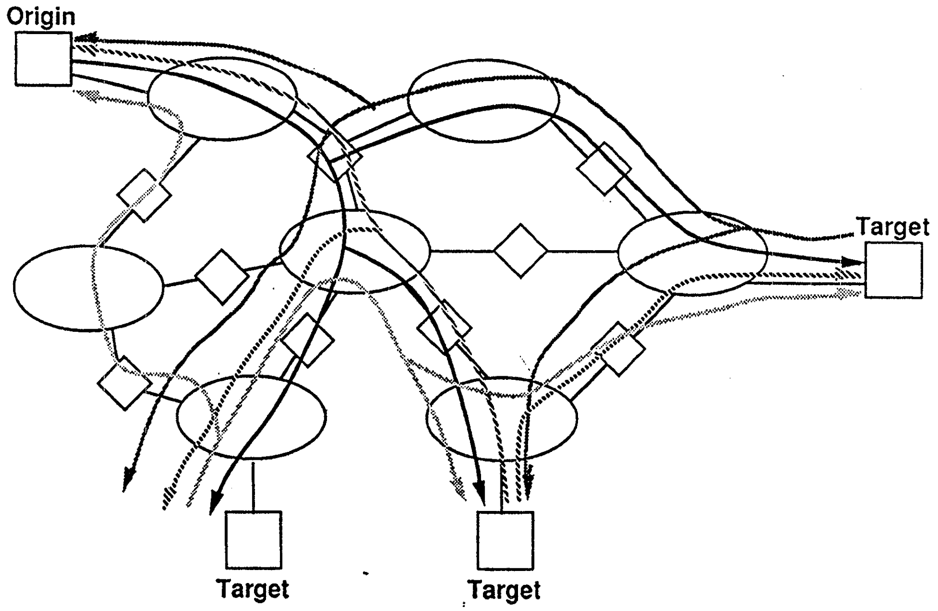

ST Features

- Separate data packet forwarding processing from stream control operations
- Minimal data processing
  - Results in potential high performance
- Simple control mechanism
- Flexibility
  - Supports arbitrary topologies
  - Allows asymmetric communication

ST Features (Cont.)

- Datagram based control protocol
- Unique Stream identifier
  - Origin IP address plus unique number
- Stream priority
  - Potentially per-target priority
- Data flow specification
- "Port" or "Socket" identifier
- No explicit point to point service
  - Optimization for 2-party streams
- Routing is a separable issue
- Subsetting of protocol

Control Protocol Stream Setup

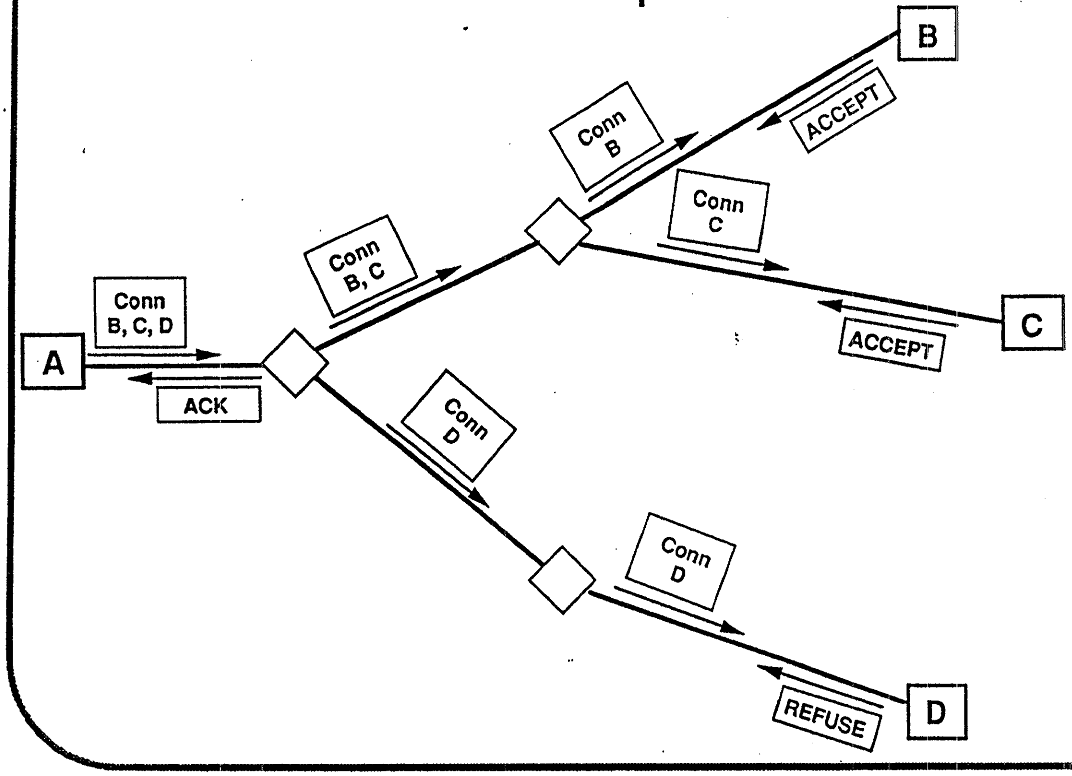

A

Conn B, C, D

ACK

Conn B, C

Conn D

Conn B

Conn C

ACCEPT

B

C

ACCEPT

Conn D

REFUSE

D

Control Protocol Stream Shutdown

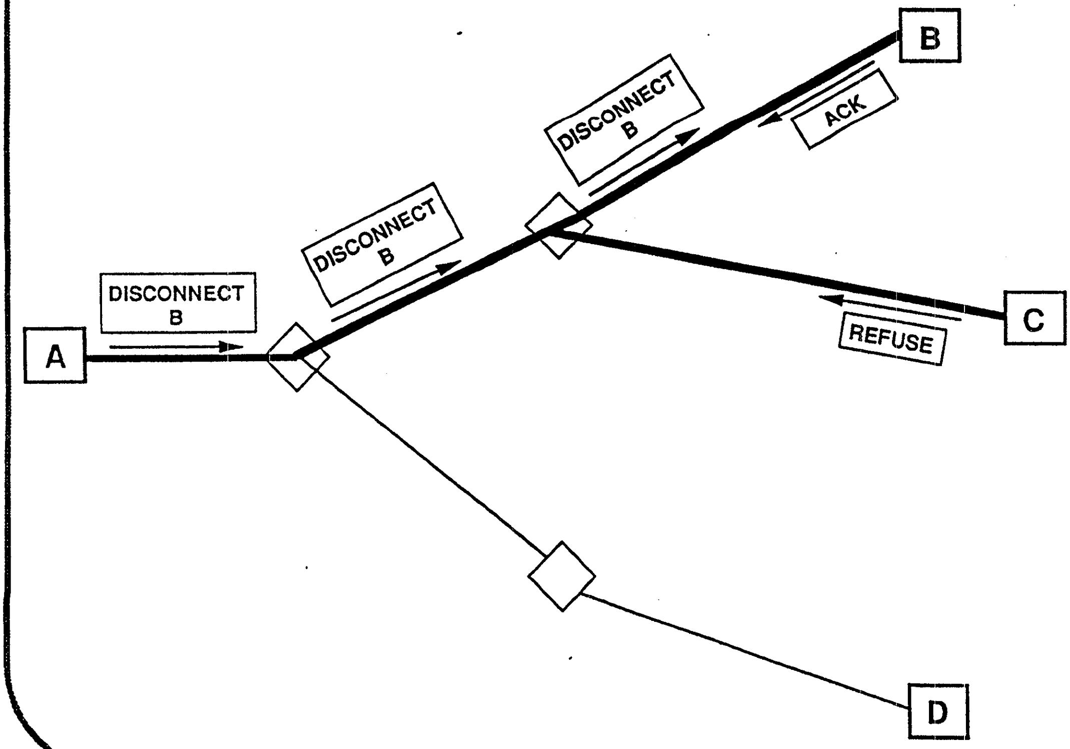

A

DISCONNECT B

DISCONNECT B

DISCONNECT B

ACK

B

C

REFUSE

D

Data Packet Forwarding

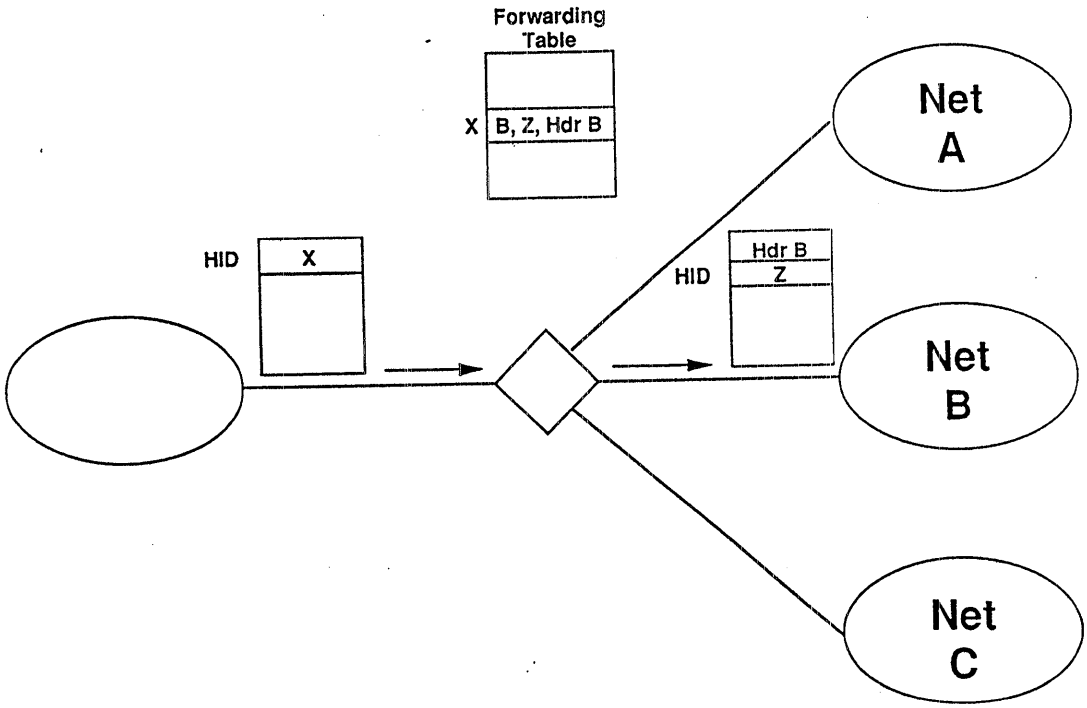

HID

- Identifies stream to which a data packet belongs
- Short, allows fast access to forwarding information
- Can change from hop to hop
- Must be unique at receiver
- Multicast network will delivery same HID to multiple next hops
- Need HID selection mechanism
  - Distributed selection
  - Previous hop selects HID
  - Next hops can reject the selection
- Current mechanism is for previous hop to select an HID randomly
  - Probably acceptable if neighbors have about 1000 streams
  - Alternately could use a centralized or distributed HID server

Data Packet Multicast Forwarding

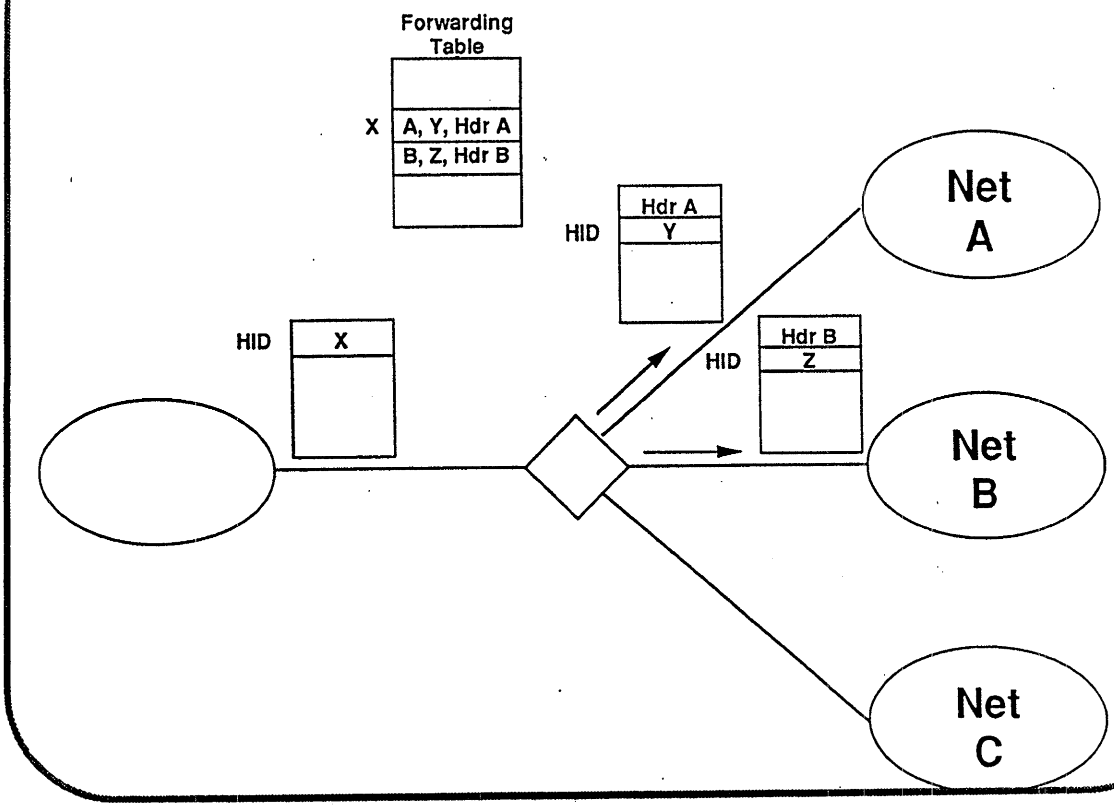

Forwarding Table

X	A, Y, Hdr A
	B, Z, Hdr B

Data Packet Forwarding Across Multicast Network

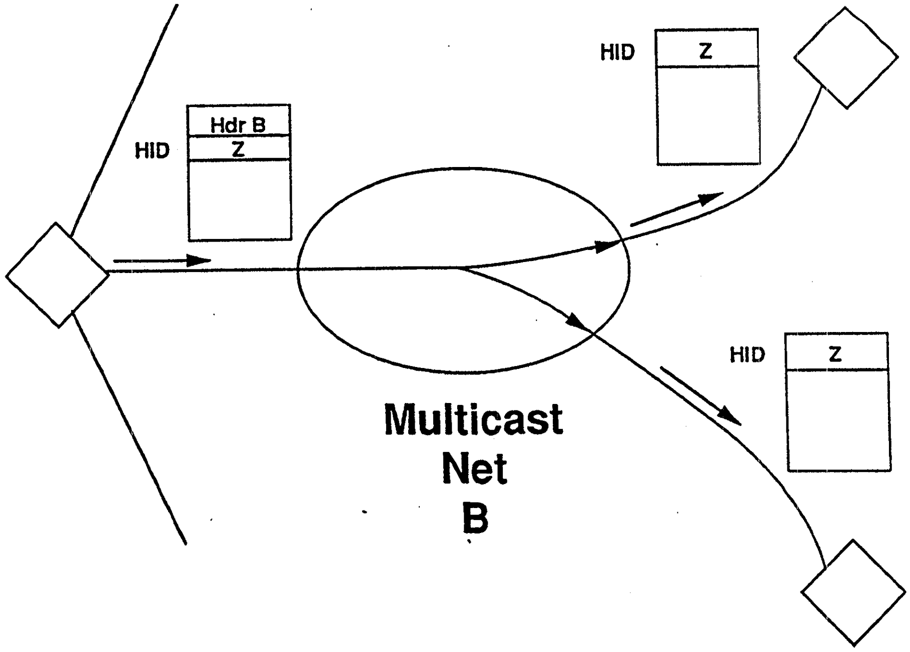

HID Collision

Problem: When adding a new branch to an existing stream the HID may not be unique at the new next hop.

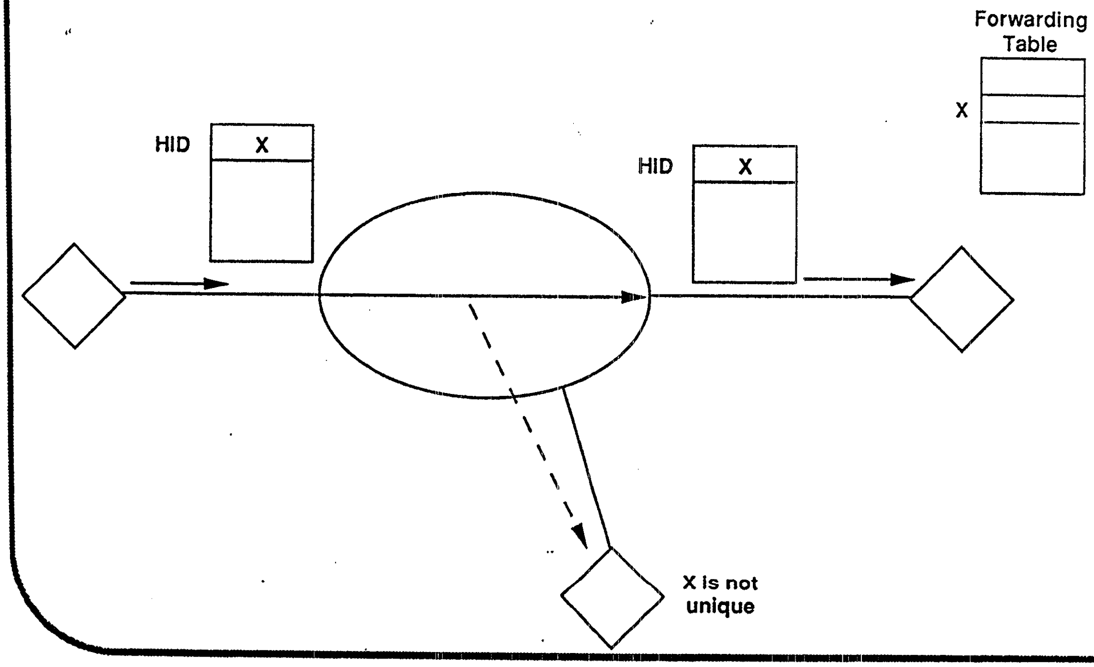

HID Collision

Solution: Allow multiple HIDs for the same stream, then reclaim the old HID.

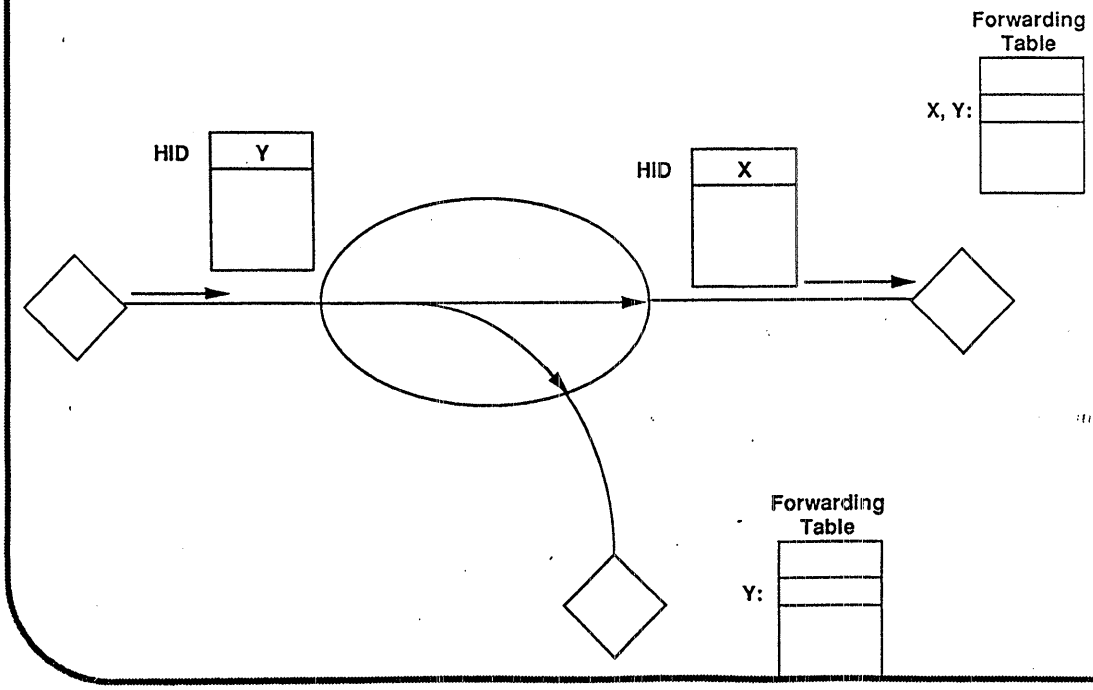

Resource Management

- Resources
  - Network bandwidth
    - currently reserve necessary bandwidth
    - try alternate routes if initial selection fails
  - Gateway buffers and CPU
    - currently over-engineer the gateways
  - Multicast network IDs
    - obtain them from the networks
- Irreconcilable conflicts result in "call" blocking or preemption

Failure Detection

(The "Me Generation" protocol)

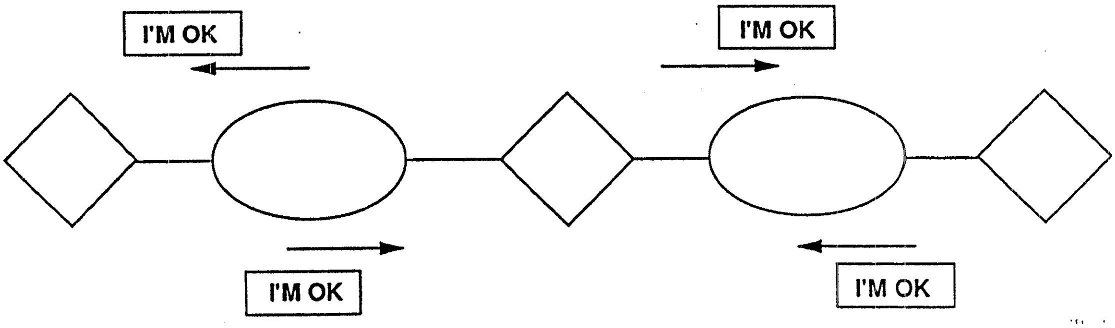

Failure Detection

- Perform handshake only between neighbors that share streams
- Neighbors exchange "lollipop" counters
- Assume that neighbor will indicated any failure of any stream it detects
- Assume that network failures will be detected either by notification by network or by interruption of this handshake
- Assume that if neighbor's counter increments properly and neighbor has not indicated a failure then all streams are still intact

Failure Repair

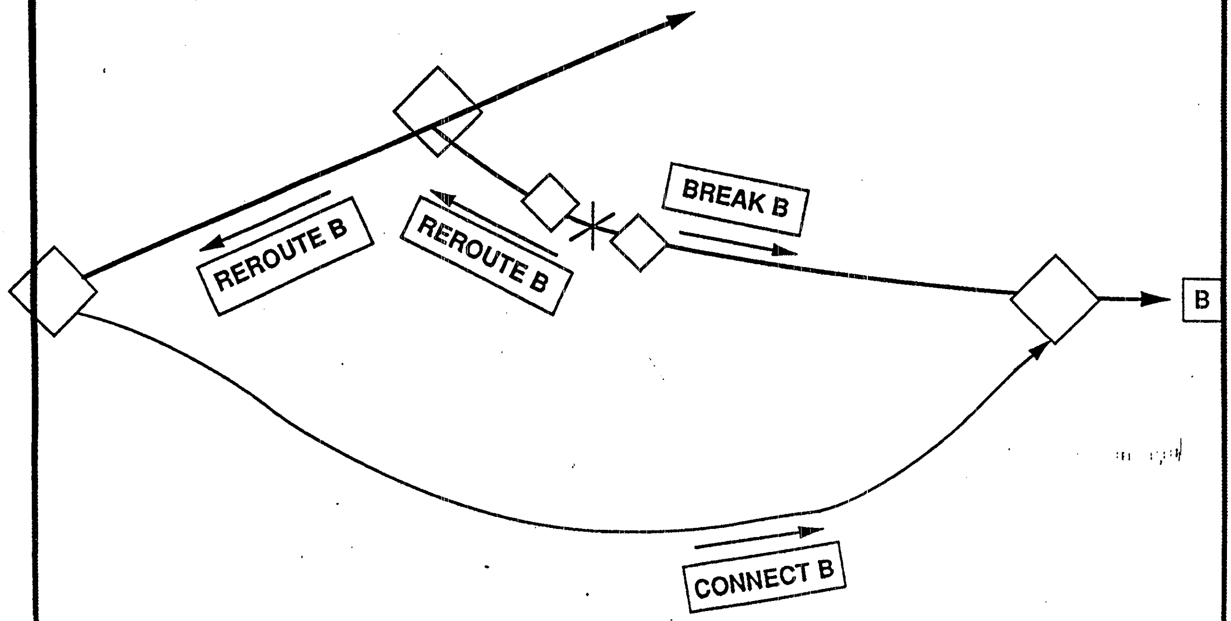

Failure Repair

- Next hop issues BREAK (similar to DISCONNECT) toward Target
- Previous hop might attempt recovery (if allowed) or issue REROUTE toward origin
- Must be careful of interaction between BREAK and new CONNECT at a downstream ST agent

Some Options

- Source route (1 target)
- "I promise to only ever have 1 target"
  - allows nets to not allocate multicast IDs
- Allow next hop to select HID
  - Makes HID selection trivial
- Build reverse path automatically
  - allows a common case to be handled with minimum packets and delays
- Groups of Streams
- IP encapsulation
- Options may restrict use of other features

BUILDING REVERSE STREAM

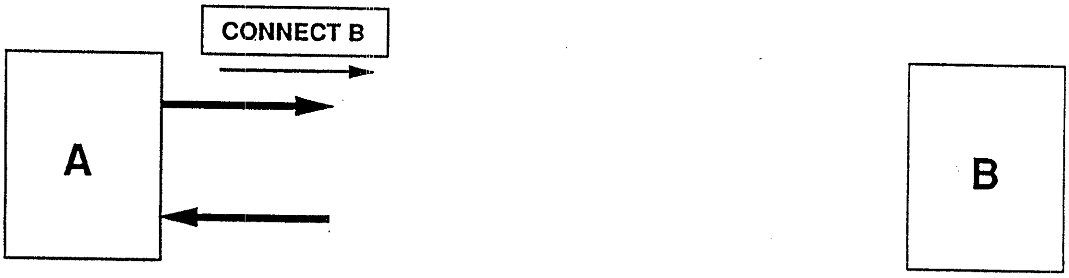
CONNECT B

A

B

BUILDING REVERSE STREAM
Continued

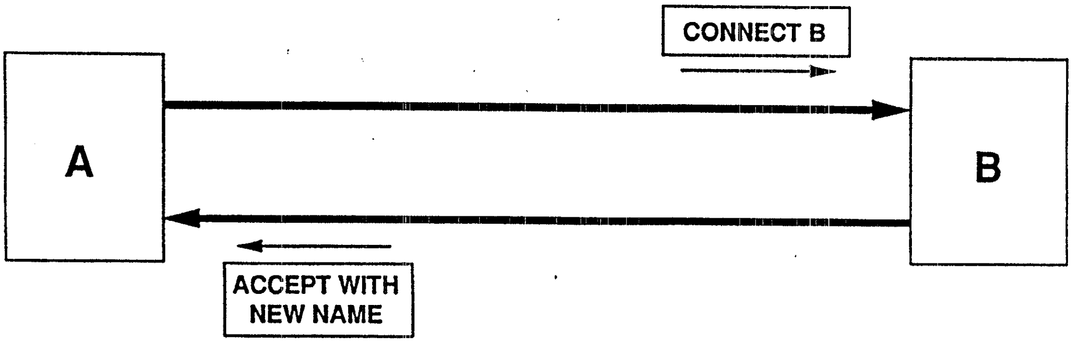
CONNECT B

A

B

ACCEPT WITH NEW NAME

Groups of Streams

- Implement shared resources
- Doesn't support "call" blocking as well as desired

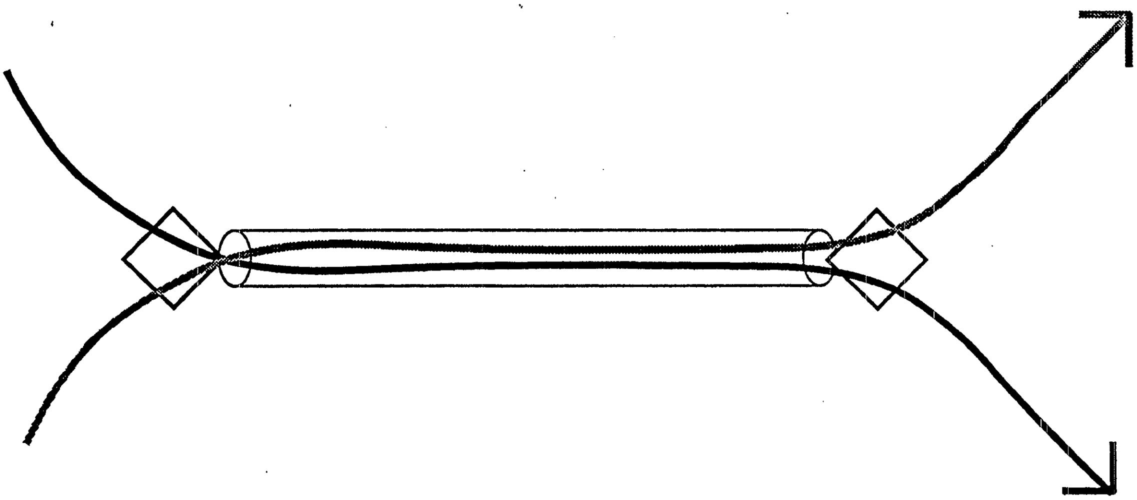

Groups of Streams
Shared Network Resources

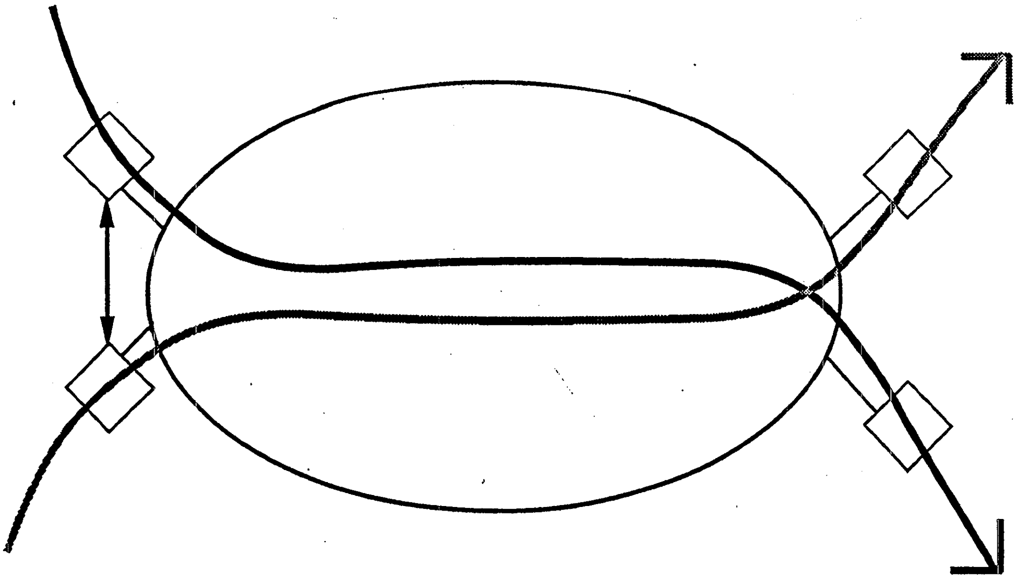

IP Encapsulation

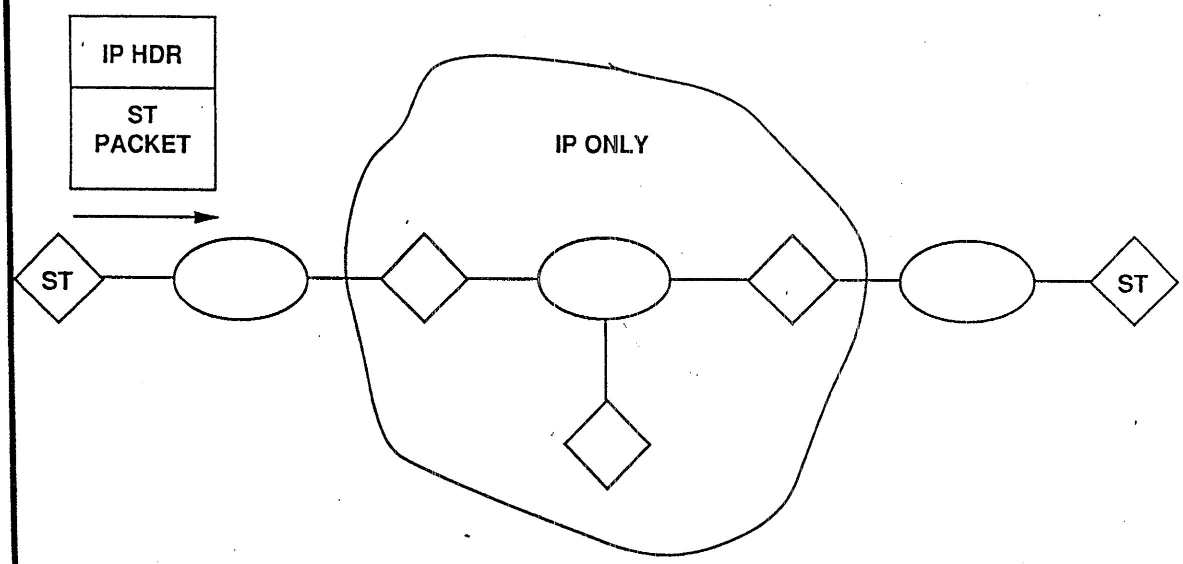
IP HDR

ST PACKET

IP ONLY

ST

ST

Other Issues

- Security with SDNS
- Routing not addressed in this document
- Monitoring of offered load is not currently implemented
- Does not directly address the communication between the ST layer and Application layer in a host

5.2 Border Gateway Protocol Status Report

Presentation by Yakoff Rekhter

The Border Gateway Protocol (BGP) is an inter-autonomous system routing protocol. It is built on experience gained with EGP as defined in RFC 904 [1] and EGP usage in the NSFNET Backbone as described in RFC 1092 [2] and RFC 1093 [3].

The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems. This network reachability information includes information on the Autonomous Systems (ASs) that traffic must transit to reach these networks. This information is sufficient to construct a graph of AS connectivity from which routing loops may be pruned and policy decisions at the AS level may be enforced.

The BGP protocol provides a high degree of control and flexibility for doing interdomain routing while enforcing policy and performance constraints and avoiding routing loops. BGP runs over a reliable transport protocol. This eliminates the need to implement explicit update fragmentation, retransmission, acknowledgement, and sequencing. Any authentication scheme used by the transport protocol may be used in addition to BGP's own authentication mechanisms. The error notification mechanism used in BGP assumes that the transport protocol supports a "graceful" close, i.e., that all outstanding data will be delivered before the connection is closed.

BGP Architecture is based on the assumption that the Internet is a collection of arbitrarily connected Autonomous Systems. The AS is assumed to be administered by a single administrative entity, at least for the purposes of representation of routing information to systems outside of the AS. BGP provides the capability of enforcing some policies based on various preferences and constraints. Policies are determined by the AS administration and are provided to BGP in the form of configuration information. These policies are enforced within a BGP speaker by affecting the selection of paths from multiple alternatives, and by controlling the redistribution of routing information. Policies are not directly encoded in the protocol.

Depending on the mechanism used to propagate BGP information within a given AS, special care must be taken to ensure consistency between BGP and the IGP, since changes in state are likely to propagate at different rates across the AS. There may be a time window between the moment when some border gateway (A) receives new BGP routing information which was originated from another border gateway (B) within the same AS, and the moment the IGP within this AS is capable of routing transit traffic to that border gateway (B). During that time window, either incorrect routing or "black holes" can occur.

In order to minimize such routing problems, border gateway (A) should not advertise a route to some exterior network X to all of its BGP neighbors in other ASs until all of the interior gateways within the AS are ready to route traffic destined to X via the correct exit border gateway (B). In other words, interior routing should converge on the proper exit gateway before advertising routes via that exit gateway to other ASs.

Fundamental to BGP usage is the rule that an AS advertises to its neighboring ASs only those routes that it uses. This rule reflects the "hop-by-hop" routing paradigm generally used by the current Internet. Note that there may be policies that cannot be supported by the "hop-by-hop" routing paradigm and which require such techniques as source routing to enforce. Discussion of whether and how these policies can be accommodated using BGP is outside the scope of this document. On the other hand, BGP can support any policy that can be accommodated within the "hop-by-hop" routing paradigm. The fact that the current Internet uses only the "hop-by-hop" routing paradigm and the fact that BGP can support any policy that can be accomplished within such a paradigm makes BGP highly applicable as an inter-AS routing protocol for the current Internet.

BGP - a pragmatic approach to the Inter-AS Routing.

Yakov Rekhter

T.J. Watson Research Center

IBM Corporation

e-mail: YAKOV at YKTVMX

April 30, 1990

The Internet Environment.

- Collection of the interconnected Autonomous Systems.
- Network Layer is IP (RFC 791).
- Transport Layer is TCP (RFC 793) or UDP (RFC 768).
- Large existing infrastructure (e.g. NSFNET).
- Limited hierarchy (by the structure of IP addresses).
- Currently used inter-AS protocol (EGP-2) is not designed for the current Internet Environment.

What is an Autonomous System ?

- Defined as part of EGP-2:
  - Set of routers under a single technical administration.
  - Single IGP and common interior metrics to handle intra-AS traffic.
  - EGP to route packets to other AS's.
- Redefined for BGP.
  - Set of routers under a single technical administration.
  - Several IGP's and several sets of metrics to handle intra-AS traffic.
  - From exterior routing an AS can be viewed as monolithic.

Inter-AS routing.

- Combination of topological and administrative constraints.
- Control over access to the AS internal resources - controlled transitivity.
- Autonomous Systems classification:
  - Transit - provides its resources to other AS's.
  - Non-transit :
    - * Multihomed - connected to more than one transit AS.
    - * Stub - connected to a single AS.

Design considerations.

- Networking architecture.
- Available technology.
- Development and deployment time frame.
- Implementation and deployment cost.
- Impact on the current infrastructure.

Requirements for the inter-AS routing.

Feasibility:

- Requirements that can not be satisfied within the design constraints.
- Requirements that can be satisfied within the design constraints, and can provide **noticable** benefits.
- Requirements that can be satisfied within the design constraints, but provide **no noticable** benefits.

Inter-AS routing components.

- Routing policies to be supported.
- Routing mechanism.
- Routing algorithms, routing protocols.
- Accounting, charging, resource allocation.
- Components are **not** independent.

On policies that can be supported within the inter-AS routing.

- Destination sensitive.
- Source sensitive.
- Path sensitive.
  - Previous hop(s) sensitive.
- Type of Service sensitive.

Routing mechanisms.

- Complete path installation.
  - End-to-end path installation.
  - Installed on all routers on the path.
  - State related to the path on all routers along the path.
- Incremental (hop-by-hop) routing.
  - Possible to have partial (e.g. hop-by-hop) path installation.
  - Some state may be present at routers.
    - * Does not reflect what route to follow.
    - * May reflect some local matters.

Complete path installation routing.

- Via connection setup:
  - Represents a significant departure from the current IP routing architecture.
  - May be impractical for large number of current applications.
  - Requires to monitor individual connections:
    - * may not scale
    - * may impose significant memory and processing requirements on routers.
- Allows complete route determination at the source.

Incremental (hop-by-hop) routing.

- Current IP routing architecture is based on incremental routing.
- No impact on any of the existing applications.
- No connections to monitor:
  - better scaling that the complete path installation routing
  - less memory and processing requirements on routers.
- Source may see the whole path; however, all the routers in the path must agree with this route.

Routing algorithms: Link State (SPF).

- Fast convergence time.
- Works well in absence of any hierarchy (but does not scale).
- Works well with controlled hierarchies (e.g. OSPF protocol) - controlled hierarchies may not be feasible for the inter-AS routing.
- Complete topological knowledge allows both path installation and hop-by-hop routing.
- Complete topological knowledge is infeasible for the inter-AS routing - complete path installation is required (alternative is to have a world-wide coordination of all routers involved in the inter-AS routing).

Routing algorithms: pure Distance Vector.

- Slow convergence time.
- Works well with arbitrary hierarchies.
- Does not require complete topological knowledge.
- Has very limited topological knowledge - next hop.
- Requires hop-by-hop routing.
- Can not support complete path installation.
- Global monotonically increased metric; globally consistent metric comparison procedures for loop suppression.

Routing algorithms: mixed (distance vector with complete path).

- Shows much better properties with respect to routing information loop suppression and convergence time than the pure distance vector algorithms.
- Works well with arbitrary hierarchies.
- Does not require complete topology knowledge, but can reconstruct the locally relevant partial topology.
- Either complete path installation or incremental (hop-by-hop) routing may be used.
- Does not require any metric.

Accounting, charging, resource allocation.

- Passive:
  - long term resource allocation
  - conforms to the current Internet architecture
  - no performance impact
  - quasi-dynamic interaction with the inter-AS routing
- Active:
  - short term resource allocation
  - significant departure from the current Internet architecture
  - potential serious negative impact on performance
  - requires complete path installation

Valid combination of choices.

Alternative 1:

- Complete path installation is mandatory.
- Link State (SPF).
- Requires controlled hierarchies.
- Supports any routing policy that can be implemented within the complete path installation routing paradigm:
  - Arbitrary destination sensitive policies.
  - Arbitrary source sensitive policies.
  - Limited set of path sensitive policies (limited by the previous hop).
  - Arbitrary Type Of Service sensitive policies.
- Routes are computed on demand.

BGP Architecture - a pragmatic approach to the inter-AS routing.

- Not "The Inter-AD Protocol For Generations To Come" -**TIADPFGTC** !!!
- Not the long term solution - long term solution is the "one that would be right for generations to come" (as defined by Professor Finnegan).
- Does not solve **all** problems for **all** people.

Policies that can be supported with BGP.

- Any policy that can be implemented with incremental (hop-by-hop) routing:
  - arbitrary destination sensitive policies
  - arbitrary path sensitive policies
  - limited set of source sensitive policies
  - limited set of Type Of Service sensitive policies
- Implemented via controlled distribution of the routing information (see also RFC1104).

Valid combination of choices (cont.).

- Alternative 2:
  - Designed for the current Internet architecture.
  - Hop-by-hop routing mechanism is mandatory, complete path installation is optional.
  - Distance vector algorithm (either pure or mixed).
  - Supports arbitrary hierarchy.
  - Supports any routing policy that can be implemented within the hop-by-hop routing paradigm:
    - * Arbitrary destination sensitive policies.
    - * Large number of source sensitive policies.
    - * Arbitrary path sensitive policies.
    - * Large number of Type Of Service sensitive policies.

BGP Architecture - a pragmatic approach to the inter-AS routing.

- Based on the current Internet architecture.
- Based on the current router technology.
- Inter-AS routing is needed *asap*.
- Minimize the implementation and deployment cost.
  - Can be implemented in less than 3 man-months on any existing commercial router.
- No impact on existing infrastructure.

Routing mechanism required by BGP.

- Incremental (hop-by-hop) routing is mandatory.
- Complete path installation is optional.

Accounting, charging, resource allocation in BGP:

- Decoupled from the inter-AS routing.
- Passive monitoring of resource utilization by means outside of BGP.
- Resource allocation by means outside of BGP (e.g. queuing).
- BGP derived information may be used for resources allocation.

BGP Algorithm.

- Mixed: distance vector with complete AS path.
- Efficient routing information loop suppression.
- Speed up convergence.
- No support for hierarchy of AS's.
  - Limited hierarchy support by IP.
- Efficient routing information loop suppression.
- Speed-up convergence.
- Rich (*but not arbitrary*) selection for Type Of Service Routing.
- Limited form multipath - restricted to the same AS path.
- Allows to detect globally inconsistent policies.

Path Attributes.

- Network reachability information has a set of attributes associated with it.
- Attributes taxonomy:
  - Well-known - recognized by all routers in all AS's.
    - Mandatory - must be present with every piece of routing information.
    - Discretionary - may not be present with every piece of routing information.
  - Optional - may not be recognized by every router.
    - Transitive - may be passed to other AS's even if not recognized (with marking as PARTIAL).
    - Non-transitive - must be dropped if not recognized.
- Extendibility of the protocol.

Inter-AS (BGP) - Intra-AS Routing interaction.

- Not all routers within an AS are running BGP.
- Potentially all the routes within an AS may carry inter-AS traffic.
- Potentially different convergence rate of Inter-AS (BGP) and Intra-AS protocols.
- Several mechanisms to ensure synchronization between BGP and an intra-AS protocol.

Syntax/Semantics for expressing policies with BGP.

- Consistent with the set of policies supported by BGP.
- Provides unambiguous way of expressing policy based routing constraints.
- Interface between legal policies (human) and actual routing (computer).

Protocol specifics.

- Incremental updates to conserve bandwidth and processing.
- Build on top of reliable transport protocol - TCP.
  - Guaranteed delivery.
  - Recover from information delivered out of order.
  - Network/Transport layer authentication/security.

What can we gain with BGP ?

- Less bandwidth and CPU utilization.
- Faster adaptation.
- Loop-free inter-AS routing.
- Rich support for policy based routing.

What can we loose with BGP ?

- EGP-2.

Current status - BGP-1.

- BGP-1 (RFC1105) was published in June 1989.
- BGP-1 is implemented by cisco, IBM, and Cornell University (*gated*).
- BGP-1 is deployed and available as a prototype in the NSFNET since October 1989.

Current status - BGP-2.

- Protocol and architecture documents were approved by the IWG during February 1990 IETF.
- Protocol document was published as an Internet Draft in February 1990.
- Protocol document passed review of the Area Director for Routing April 12, 1990 - *"BGP be published as an RFC and entered in the standard track as a Proposed Standard after some editorial changes and clarifications"*.
- Protocol document was presented to the Area Director for Routing with requested editorial changes and clarifications April 23, 1990.
- Architecture document was published as an Internet Draft in March 1990 and given to the Area Director for Routing.

Chapter 6

Technical Presentations

6.1 The Pennsylvania Research and Education Network

Presentation by Tom Bajzek/ PREPNET and Walt Burmeister/ Bell of PA

Walt Burmeister/ Bell of PA

PREPnet was created when a peer group of Pennsylvania research university administrators from seven of the major research universities approached Bell and asked about the feasibility of developing a mid-level network in Pennsylvania.

These requests were prompted by a number of needs:

- sharing resources;
- facilitating scholarly collaboration;
- accessing supercomputers and other data bases;
- and easy, economical access to the Internet.

Also, the emergence of other regional networks threatening to encroach into Pennsylvania spurred this group to ask for a Pennsylvania-based network – a network that would respond to the needs of researchers, educators, government and businesses in Pennsylvania.

These forward-looking Vice Presidents of Computing and Research were also responding to discussions with the National Science Foundation and other higher education research communities.

Then – and now – they are all working to develop a National Research and Education Network, the NREN.

It was clear that in Pennsylvania, a mid-level network was an essential factor in positioning the Commonwealth for connectivity to this larger national effort to share data and information.

In order to respond to this request to provide connectivity across Pennsylvania, we had to find a partner. With divestiture from AT&T, Bell of Pennsylvania can no longer provide long-distance services. We can no longer provide direct access across the state. We are limited to providing service within six regional areas.

So we enlisted the support of the Commonwealth of Pennsylvania. They have a network crisscrossing the state, and they agreed to provide spare capacity on their T-1 circuits to serve as the backbone for this new network.

The Commonwealth made a natural partner – an opportunity to meet and match needs. They could provide long-distance connectivity and spare capacity. In exchange, the state would accrue the benefits of an infrastructure that would serve to improve economic growth and development.

I will expand more on the economic development shortly. Let me return to a few more specifics about the network.

To provide ready access for our founding universities, we built two hubs in our central offices – one located in Philadelphia and one in Pittsburgh. Initially, these two hubs connected the seven founding universities:

- Carnegie Mellon
- The University of Pittsburgh
- Penn State University
- The University of Pennsylvania
- Drexel University
- Temple University
- and Lehigh University.

To provide access to the Internet community, the Pittsburgh Supercomputing Center was enlisted to serve as PREPnet's gateway. Besides operating and maintaining PREPnet, Bell also agreed to promote and market this service. We have committed substantial resources to this investment.

Why did we, Bell of PA, become involved with developing this network? Why are we the only local operating company investing resources to develop a mid-level network? The answer lies in the return on our investment. We believe that PREPneet is a new initiztive - and like any investment, it brings both risk and reward.

What do we expect for our return on this investment? We expect economic growth and development for Pennsylvania.

We are a company committed to a future in Pennsylvania. We serve customers in Pennsylvania. Our vision of the future is dependent on continued economic development in this state. To paraphrase an old saying, as Pennsylvania goes, so goes the state of Bell of PA. By providing the electronic infrastructure necessary for the research and education community, we are bringing Pennsylvania abreast of other states who already have similar networks.

PREPnet provides researchers with the opportunity to improve productivity by reducing access time to information. PREPnet can make a specialized scientific laboratory become a virtual laboratory by supporting visualization, on-line collaboration, and real-time control. PREPnet can provide faster peer interactivity and review...independent of distance.

The value of information is determined by the user. With PREPnet, we can provide the means for accessing and sharing research and information. As information is shared, it is leveraged and grows, and the developing synergy is totally independent of location and distance.

The impact of research lies not only in collaboration and sharing, but also in technology transfer; in transferring research results from the educational sector to the business sector, and within the business sector, from one business to another and eventually it is translated into consumer products to enhance the state's economic development.

This PREPnet partnership is providing a critical tool necessary for Pennsylvania to compete not only nationally, but internationally. So, a major priority – economic development – is both a need and a benefit.

Another reason for Bell's participation is the opportunity to gain experience working with protocol-dependent networks such as PREPnet.

TCP/IP, as this audience well knows, is a protocol that has a vital role in today and tomorrow's networking. Bell wanted to expand our knowledge about it. Again, it is the leveraging of information. What we learn about TCP/IP within PREPnet, we are applying to applications for other products and services that we are developing.

We are also gaining experience operating a mid-level network. We have experience with central office-based local area networks. But as I mentioned earlier, given today's constraints of the Justice Department, we cannot be the sole provider of statewide networks.

We would like to see the regulations lifted. But for today, PREPnet is a creative initiative to help meet customer's needs. It is an excellent opportunity to partner with others and provide connectivity and reliability in a different arena.

In addition, we, the PREPnet partnership, want to serve a larger goal. We are creating a network that can respond to the needs of different types of constituents across the state. We do not want a network just for the exclusive use of the wealthier, larger and more prestigious universities, nor do we want to limit the network to just the educational community. Our goal is to make it available to all sectors, regardless of size or geographical location. That makes PREPnet different from a commercial network.

PREPnet is chartered as a non-profit organization. PREPnet has a price structure for members that provides a choice of pricing options based on the needs of the enterprise, its size and type. The PREPnet steering committee has priced the membership fees low.

We do not want to create barriers to entry for small users, whether they are colleges, government agencies, or researchers in private industry. Many others can now access the same wide range of resources previously available only to large universities – and they can do it at a choice of different transmission speeds and access options – depending on that their needs and budgets are.

Since PREPnet's inception in 1988, we have built two additional hubs: one in Scranton and another in Allentown. These hubs are part of our response to the growing demands of customers who need economical access across the state. We have also provided a terminal server in the Philadelphia hub to accommodate slow speed asynchronous access.

In addition to managing and operating this network, Bell of PA serves as the marketing arm of the partnership. Today, we have 27 PREPnet subscribers, ranging from higher education to steel manufacturing, from software development to medical research.

We are continuing to move forward to achieve PREPnet's four major goals:

- Contribute to the state's economic development

- Facilitate the sharing of research information
- Increase the opportunities for collaborative research
- Enhance the competitive posture of the state.

As demand increases and resources are available, the PREPnet partnership plans to build several more hubs. We may eventually have to provide a bigger pipeline. While T-1 is fine for the present, as the critical mass of users and their applications grow, PREPnet will also grow.

Right now, Bell of PA is working with CMU and the Pittsburgh Supercomputing Center for the provision of a three gigabit line.

In Philadelphia, we are developing a metropolitan area network for Temple University – probably one of the first of its kind.

I see the potential for expanding the functionality of PREPnet through these technologies.

All of us here today are involved in the technology explosion that is occurring. In many cases, technology seems to be growing as fast or faster than we can apply it.

In today's technological world no single resource can provide all the information, technology and resources needed to compete effectively in the global marketplace.

It is a bold, creative joint initiative that combines partners and technology to provide a network that can change the way in which learning and research occur.

PREPnet accelerates the flow of research to other researchers and to commercial enterprises. Bell of PA has invested in PREPnet because we believe it will significantly impact the economic growth within Pennsylvania. We believe that this network solution provides a partnership from which all of us will gain more than we have invested.

PREPnet Map

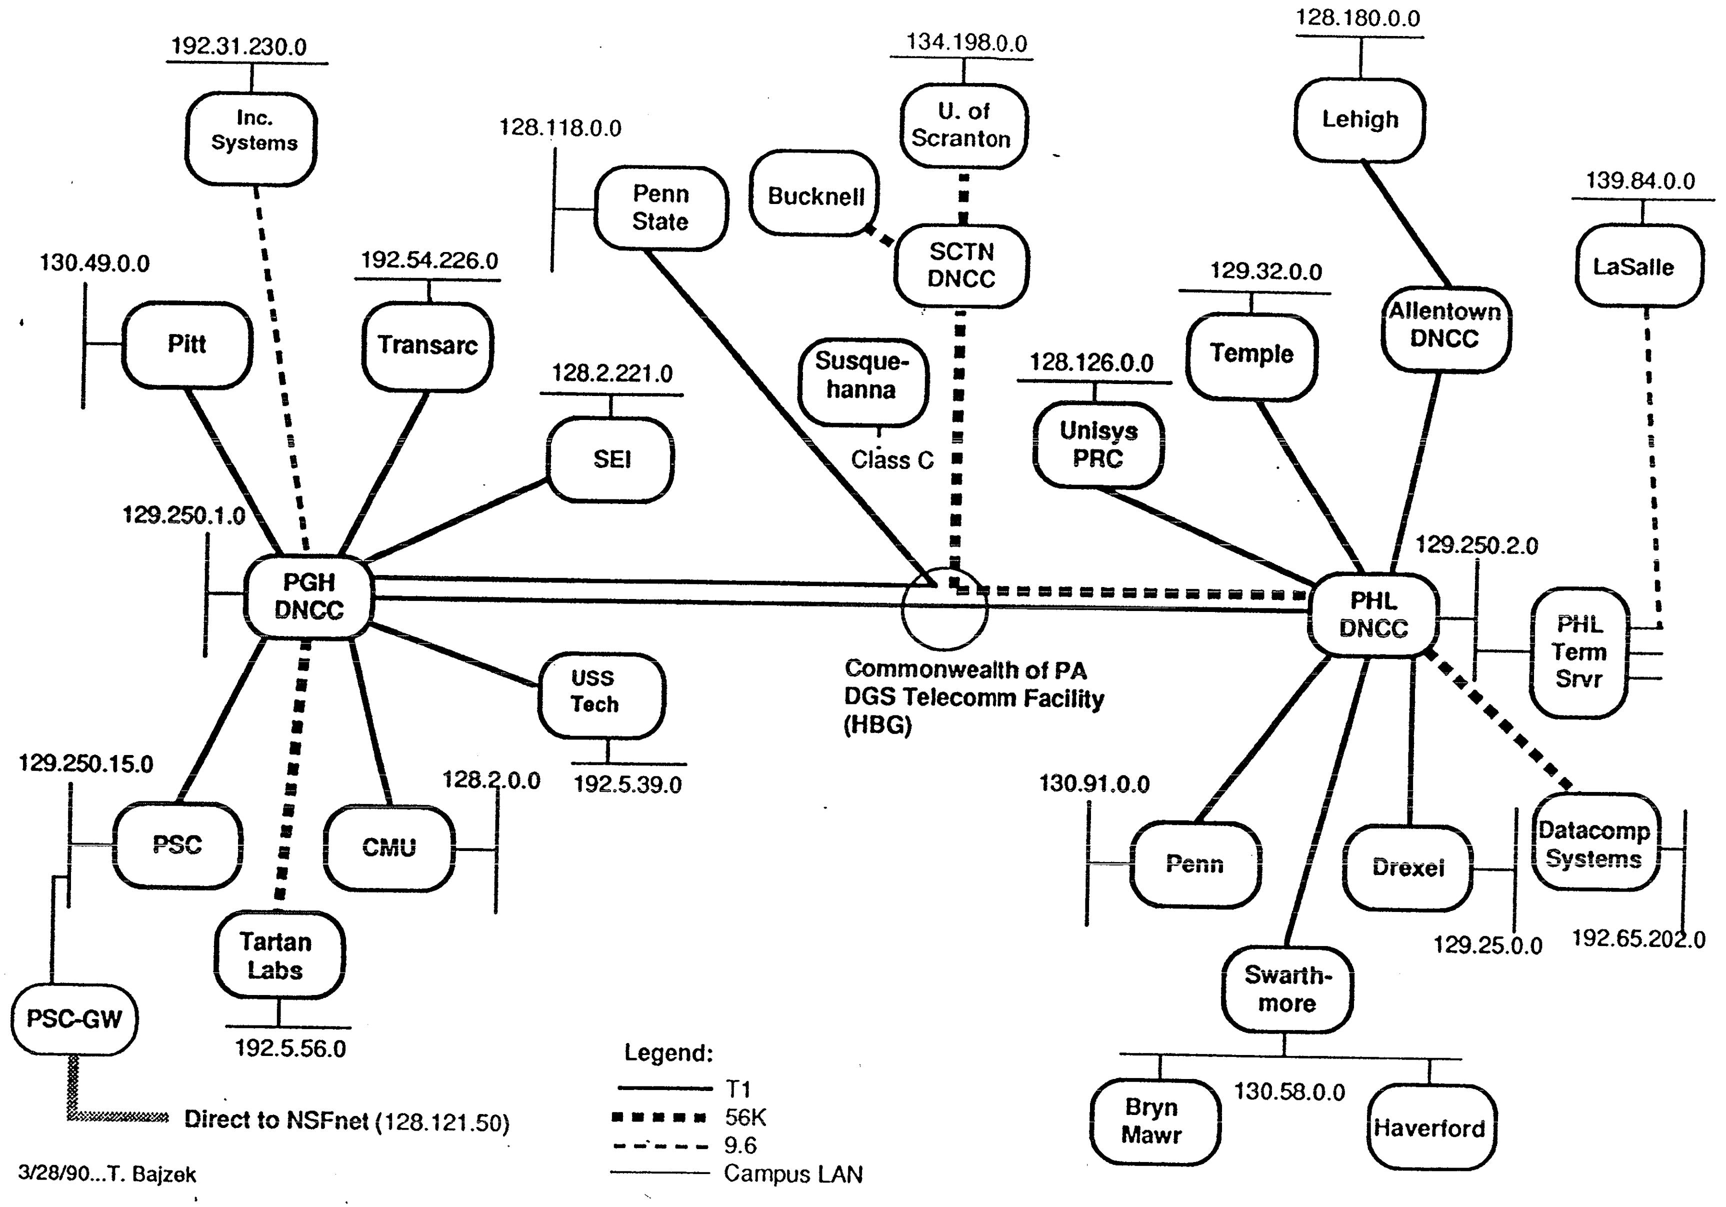

3/28/90...T. Bajzek

PREPnet Membership

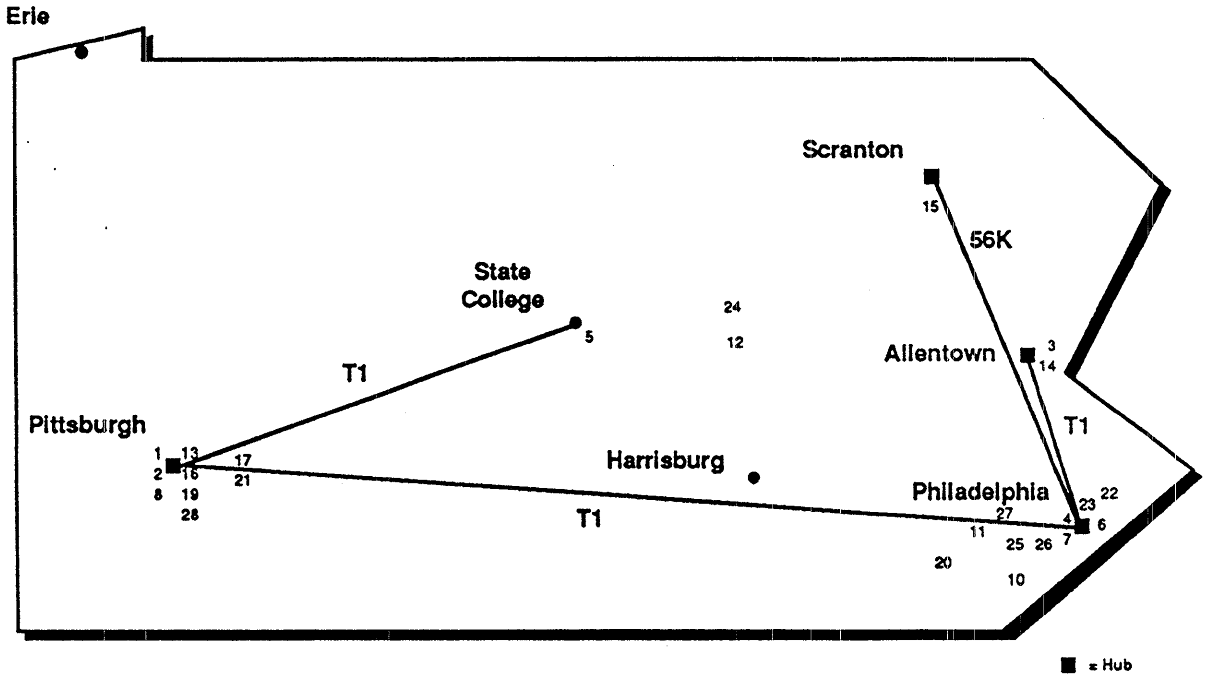

1 Carnegie Mellon University
2 Universityof Pittsburgh
3 Lehigh University
4 University of Pennsylvania
5 Penn State University
6 Temple University
7 Drexel University
8 Pittsburgh Supercomputing Center
9 (Reserved)
10 Swarthmore College
11 Unisys Paoli Research Center
12 Susquehanna University
13 Transarc Corporation
14 NET Ben Franklin Tech Center
15 University of Scranton
16 Software Engineering Institute
17 USS Technical Center, USX Corporation
18 Incremental Systems Corporation
19 Allegheny-Singer Research Institute
20 West Chester University
21 Tartan Laboratories, Inc.
22 Datacomp Systems, Inc.
23 LaSalle University
24 Bucknell University
25 Bryn Mawr College
26 Haverford College
27 Soft Switch, Inc.
28 Maya Design Group

PREPnet Information Services

Provider/Service	Description
Provider: CMU Service: LIS Availability: **now** IP-address: cmulibrary.andrew.cmu.edu Requirements:vt100, etc.	LIS is a library catalog information retrieval system which indexes all words in the entries. Several databases are available: the CMU library catalog and journal list, a set of bibliographies on a variety of topics compiled by CMU librarians, and in index to architectural pictures found in a number of standard reference books.
Provider: Penn State Service: LIAS Availability: **now** IP-address: lias.psu.edu Requirements: vt100	LIAS is the on-line catalog of the Penn State libraries.
Provider: Penn State Service: PENpages Availability: **now** IP-address: psupen.psu.edu Requirements: vt100 User: pprepnet	PENpages is a database of agricultural and Extension related information ranging from daily, weekly, and monthly agricultural news and alerts to permanent reference material. It is maintained by the faculty and staff of the College of Agriculture and offered as a public service by the Cooperative Extension of Penn State.
Provider: Penn State Service: EBB Availability: **now** IP-address: cac3270.psu.edu Requirements: vt100 IP-address: psuvm.psu.edu Requirements: tn3270	EBB is a database of information relating to Penn State, including academic programs, academic calendars, and phone directories.
Provider: Penn State Service: EDIN Availability: **now** IP-address: cac3270.psu.edu Requirements: vt100 IP-address: psuvm.psu.edu Requirements: tn3270	The Pennsylvania State Data Center maintains this database of population and economic statistical data, which includes, among other things, the Commerce Business Daily. EDIN is accessible through the EBB service of Penn State.
Provider: Lehigh U. Service: ASA Availability: **now** IP-address: asa.lib.lehigh.edu Requirements: vt100	Lehigh's (GEAC-based) library catalog is now accessible through PREPnet.
Provider: U. of Pennsylvania Service: PennLIN Availability: **now** IP-address: pennlib.upenn.edu Requirements: vt100	PennLIN is Penn's NOTIS-based on-line library catalog system.

PREPnet Information Services

Provider: U. of Pennsylvania
Service: MEDINFO
Availability: **now**
IP-address: med.upenn.edu
Requirements: vt100
User: penn_med

UPenn's medical school provides access to a version of its MEDINFO bulletin board via PREPnet.

Provider: U. of Pittsburgh
Service: NIH Guide
Availability: **now**
IP-address: nic.cis.pitt.edu
Requirements: vt100

The NIH (National Institutes of Health) Guide contains information on scientific initiatives and administration regarding extramural programs. The guide is being provided to Pitt on-line as a pilot project. Files are available by anonymous FTP in the "nihguide" subdirectory. The Guide is published weekly, and information from the last 4 weeks will be online.

Provider: U. of Pittsburgh
Service: PITTCAT
Availability: July 1, 1990
IP-address:
Requirements: vt100

PITTCAT is Pitt's NOTIS-based on-line library catalog system.

Provider: U. of Pittsburgh
Service: NIAC
Availability: **now**
IP-address: nic.cis.pitt.edu
Requirements: vt100

NIAC, the NASA Industrial Applications Center at Pitt provides the Federal Laboratory Directory. The database contains information pertaining to research centers, facilities and laboratories which function under the direction of the US government. Files are available by anonymous FTP in the "niac" subdirectory.

Provider: Temple U.
Service: Aff. Action Planning
Availability:
IP-address:
Requirements:

This database contains data on the impact and benefits of affirmative action on a major hospital construction project. It appears to be a case study of an actual project.

Provider: Temple
Service: Wage & Empl. Data
Availability:
IP-address:
Requirements:

One of these databases contains employment and salary data on scientists and engineers in a variety of fields, and the other traces wage and employment growth in a variety of industries in Pennsylvania.

Provider: Temple
Service: Library Catalog
Availability:
IP-address:
Requirements:

Temple's library catalog could be available on-line, along with a personalized information search and document retrieval service.

6.2 Bringing X.400 to the Internet

Presentation by: Allan Cargille

The University of Wisconsin is funded by NASA, on behalf of the FRICC, to conduct a pilot project to bring X.400 to the Internet. We are running X.400 over ASN.1, Session, RFC1006, and TCP/IP. Our X.400 software was written as part of the IBM-funded Wisconsin ARGO project. IBM is allowing free distribution to twenty college, university, government and non-profit research centers.

The software supports a modified ucb/mail user interface. Mail files are stored in ASN.1 "Abstract Syntax Notation," a binary format. X.400 address formats (O/R Names) are used. Additional features supported include deferred delivery, delivery notification, test messages (probes), and multiple body parts. Currently only ASCII text is supported.

At present, the software is running on approximately nine machines at the University of Wisconsin and NSF. We are in the process of getting license agreements signed and installing the software at a number of other sites, including NASA-Ames, Rice University, Harvard University, the University of Pennsylvania, and the University of Sydney. Several more of these sites should be online by early May.

In the initial phase, we have created an isolated island of X.400 electronic mail. Initial users can only mail to other members of our project. There are no bridges to international X.400 networks or to the Internet mail world. Our first step toward connectivity is to establish connections to the European RARE X.400 project. This will establish connectivity with an estimated 12,000 person European academic and research community. Currently we are testing X.400/ASN.1/Session/RFC1006/TCP/IP connections to France and Norway. The second step, integrating X.400 mail with standard Internet mail, will be accomplished by operating an RFC987 gateway at the UW. The public domain X.400 product "PP" from the University College London contains an RFC987 gateway. The UW is a beta-test site for PP.

X.400 naming issues present complex challenges that have not been resolved yet for the United States. What structure the U.S. portion of the Internet will take is still an open question. One possibility is that the Internet will become a Private Management Domain (PRMD). Rather than wait for all of these issues to be resolved for the U.S., we have chosen to operate an experimental PRMD named the "XNREN." We are acting as a naming authority for organizations within our PRMD. When official decisions have been reached about the role of the Internet in X.400 naming, we will conform to those decisions and merge our namespace into the new structure.

As PP becomes publicly available, we are planning on incorporating more sites into the experiment. We would like to exchange addressing information and establish connectivity with any Internet parties running X.400, and distribute Internet addressing information to all interested sites.

Bringing X.400 to the Internet

May 1, 1990

Allan Cargille
Manolis Tsangaris
`Larry Landweber

University of Wisconsin

cargille@cs.wisc.edu

/c=us/admd=" "/prmd=nren/
/o=uw-madison/ou=cs/s=Cargille/g=Allan/

Features of X.400

- Separation of delivery information and contents
- Message transferred as data structure encoded in binary
- Multiple body parts
- Different Body Parts
- International Standard
- Improved naming:

  /c=us/admd=" "/prmd=nren/
  /o=uw-madison/ou=cs/s=Cargille/g=Allan/

New Services

- Well-defined delivery report
- Probe (test message)
- Delayed Delivery
- Automatic translation of body parts
- Others

X.400 Model

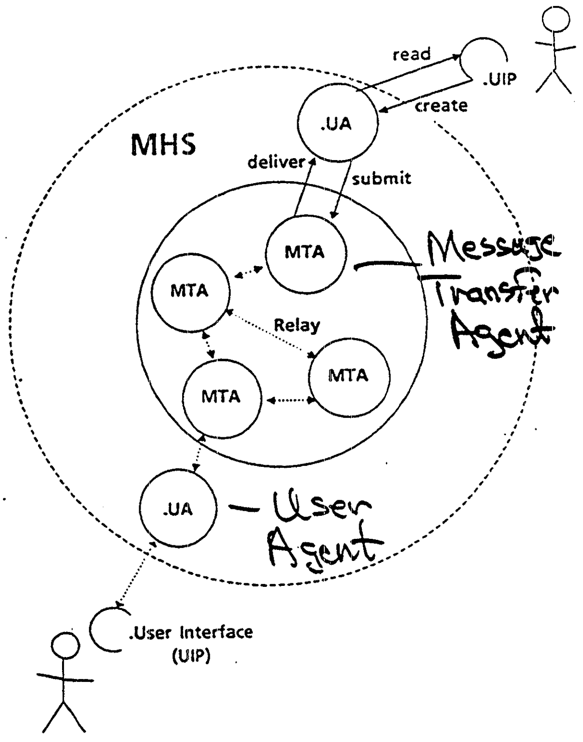

Message Transfer Agent

User Agent

M12

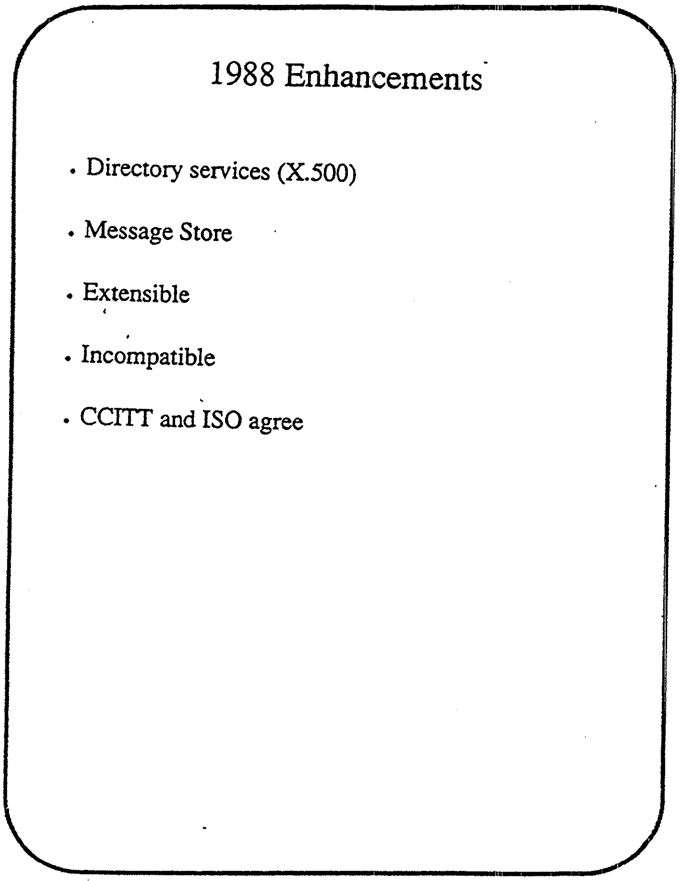

1988 Enhancements

- Directory services (X.500)
- Message Store
- Extensible
- Incompatible
- CCITT and ISO agree

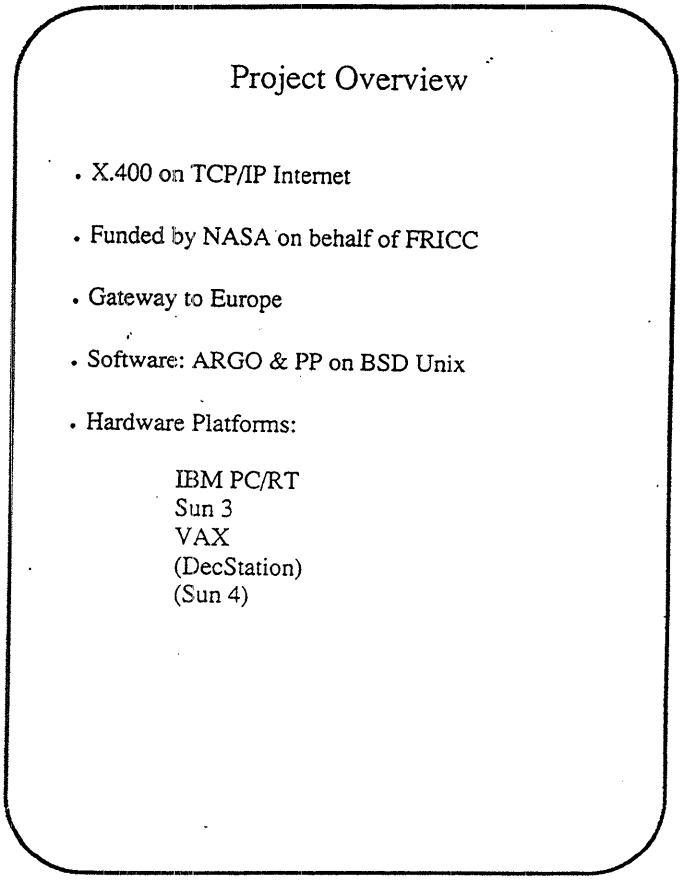

Project Overview

- X.400 on TCP/IP Internet
- Funded by NASA on behalf of FRICC
- Gateway to Europe
- Software: ARGO & PP on BSD Unix
- Hardware Platforms:

  IBM PC/RT
  Sun 3
  VAX
  (DecStation)
  (Sun 4)

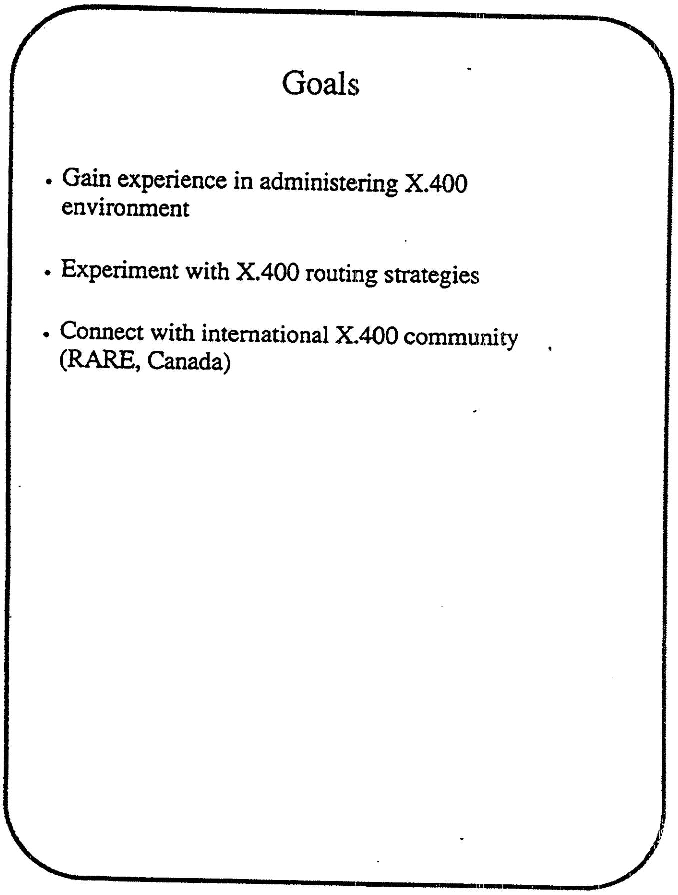

Goals

- Gain experience in administering X.400 environment
- Experiment with X.400 routing strategies
- Connect with international X.400 community (RARE, Canada)

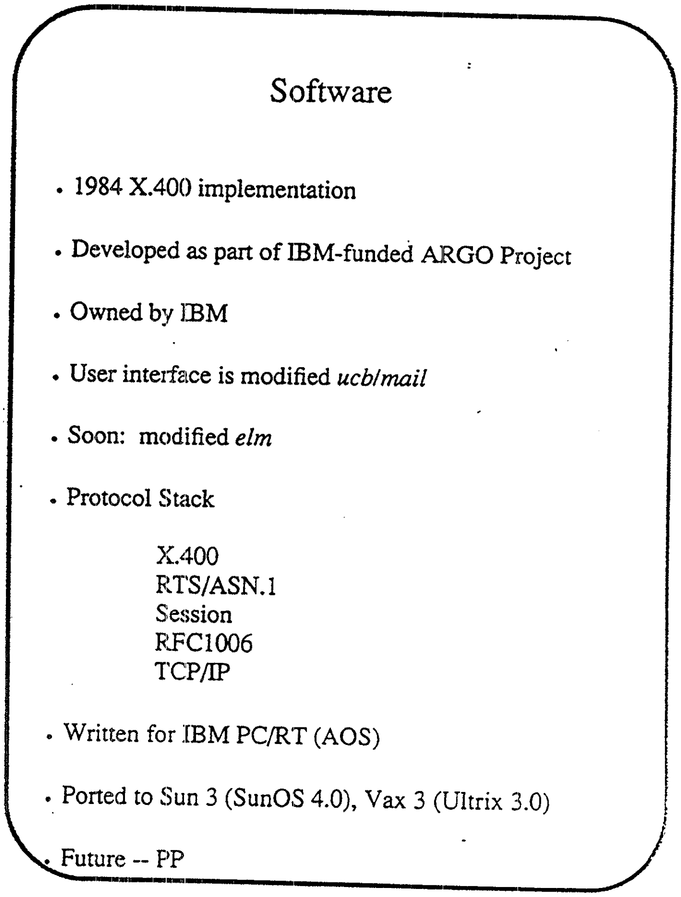

Software

- 1984 X.400 implementation
- Developed as part of IBM-funded ARGO Project
- Owned by IBM
- User interface is modified *ucb/mail*
- Soon: modified *elm*
- Protocol Stack

  X.400
  RTS/ASN.1
  Session
  RFC1006
  TCP/IP

- Written for IBM PC/RT (AOS)
- Ported to Sun 3 (SunOS 4.0), Vax 3 (Ultrix 3.0)
- Future -- PP

Participants

NSF
UW-Madison
Rice University

MITRE Networking Center
University of Pennsylvania

MIT
NASA
Harvard University
University of Maryland
Merit Corporation (NFSNET)
Univ. of California, Irvine
National Library of Medicine
University of Pittsburgh Medical Center
Defense Communications Engineering Center
(Canada, Australia, Netherlands)

- Seeking additional sites

International X.400 Activity

- Europe: RARE

  * Started in 1987
  * Most European countries
  * Well-integrated
  * over 400 sites
  * Over 15,000 users
  * Variety of implementations (15+)

- Canada
- Korea
- Australia
- U.S. way behind

International Connections

- Develop U.S. WEP
- Use international TCP/IP lines
- Connections to France and Norway
- Relay sites
- Good test of OSI
- Working on other connections

Technical Issues

- Mapping X.400 to existing structures

  * Administrative Mgmt Domains
  * Private Mgmt Domains
  * Role of Internet

- Naming issues

  * Naming authority(ies)
  * U.S. naming ambiguous

- Routing
- Bilateral agreements
- Operating Private Mgmt Domain

1990 Activities

- More participants
- Develop NREN PRMD
- Incorporate PP into experiment
- Operate X.400-RFC822 gateway
- More European connections
- More interoperability

Future Technical Activities

- Migration to X.400(1988)
- Integration with X.500 directory service
- Use over TP4/CLNP
- X.400 security extensions
- Multimedia extensions
- DNS extension to support 987/1138 mappings

6.3 The Knowbot Information Service

The loose confederation of networks that exchange electronic mail, including the Internet, CSNET, BITNET and UUCP provides services to hundreds of thousands of users. Locating a specific user in this collection of networks is a difficult problem; there is no single directory service in which all users are represented. Searching for a user requires that queries be directed to a collection of more-or-less well-known directory services, including `whois@NIC.DDN.MIL`, `whois@SH.CS.NET`, Profile and X.500. Of course, each of these services accepts different styles of queries and returns its responses in a unique format.

While the White Pages Pilot Project is exploring the use of X.500 as the basis for an Internet-wide directory service, we believe that the current diversity of directory services will persist for some time. There are also users of other networks not a part of the Internet that are unlikely to become part of any Internet-based directory service. The likelihood that users will want to query multiple directory services through a single, uniform interface motivates the development of our new tool.

We have developed the Knowbot Information Service (KIS) that uses Knowbots to automate the process of searching multiple directory services. KIS is a directory service user interface that gives the appearance of integrating heterogeneous directory services into a single, uniform service. In response to user input, KIS forwards queries to multiple directory services. The responses are collected and processed by KIS into a standard format for presentation to the user. The design of KIS is flexible and extensible, allowing convenient integration of new directory services into KIS as they become available.

Knowbots

Knowbots are programs that know about information resources and how to interact with them to obtain answers to questions

Knowbots...

- are mobile
- persist across system incarnations
- interact with other Knowbots
- reproduce when necessary

The Knowbot Information Service is a testbed application for the study of Knowbots

Locating User Information in the Internet

- User must know where to look
- User must know how to ask
- User must know how to read results

Example

Consider finding information about Tim Korb:

```
% whois -h nic.ddn.mil korb
     Korb, John T. (JTK1)            JTK@CS.PURDUE.EDU
        Purdue University
        Department of Computer Science
        West Lafayette, IN 47907
        (317) 463-3644

        Record last updated on 24-Aug-89.
% whois -h sh.cs.net korb
     There is 1 match to your query:
     -----------------------------
     Name:
             John T. Korb
     Account:
             jtk,purdue,purdue
     Ident:
             1471
     Updated:
             Thu Jan 26 17:43:40 1989
     Mailbox:
             jtk@cs.purdue.edu
     Phone:
             317-494-6184
     Address:
             Department of Computer Science
             Purdue University
             West Lafayette, IN 47907
     Misc:
             Director of Research Facilities
             systems, networks, windows
             Tim
```

The Knowbot Information Service

...provides uniform access to Internet white pages services

- Maintains an internal list of known directory services
- Accepts uniform query syntax
- Returns uniform response format

Using the Knowbot Information Service

```
% netaddress korb
              Knowbot Information Service
              ------- ----------- -------
                      Version 1.0
Copyright Corporation for National Research Initiatives 1989

Name:         John T. Korb
Organization: Purdue University
Address:      Department of Computer Science
City:         West Lafayette
State:        IN
Country:      US
Zip:          47907
Phone:        (317) 463-3644
E-Mail:       JTK@CS.PURDUE.EDU
Source:       whois@nic.ddn.mil
Ident:        JTK1
Last Updated: 24-Aug-89

Name:         John T. Korb
Organization: Department of Computer Science
Address:      Purdue University
City:         West Lafayette
State:        IN
Country:      US
Zip:          47907
Phone:        317-494-6184
E-Mail:       jtk@cs.purdue.edu
Source:       whois@cs.net.sh
Ident:        1471
Last Updated: Thu Jan 26 17:43:40 1989
```

Using the Knowbot Information Service (Continued)

```
No matches found for korb from service profile@nri.reston.va.u

Name:         Tim Korb
Address:      CS-208
Phone:        463-3644
E-Mail:       jtk@purdue.edu
Source:       profile@gwen.cs.purdue.edu
Last Updated: (unknown)

No matches found for korb from service mitwp@mit.edu
No matches found for korb from service mcimail@nri.reston.va.u
```

`netaddress` Interface

- `whois`-like syntax (command line invocation)
- Or, user can form and issue queries through interactive user interface
- User can specify:
  Name of target
  Services to search
  Service-specific identifier
  Organization
- Default services: `whois@nic.ddn.mil`, `whois@sh.cs.net`, Profile, `finger@mit.edu`, MCImail
- Other services: NYSERNet White Pages Pilot Project, `finger@UNIX_host`

Knowbot Information Service Architecture

- `netaddress` client provides user interface and sends queries to Knowbot Information Service server
  - Currently ASCII-only interface
  - X interface ''soon''
- Server reformats and forwards queries to Internet white pages services
- Server fields replies, restructures into standard format and returns to `netaddress` client

Internals

- netaddress client and the Knowbot Information Service server communicate through TCP
- Server uses awk to translate queries and replies
- New services can be incorporated by adding new awk scripts

What's Behind netaddress

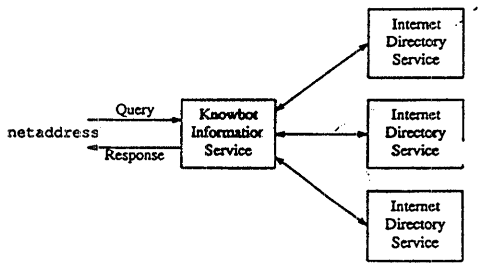
netaddress
Query
Response
Knowbot Informatior Service
Internet Directory Service
Internet Directory Service
Internet Directory Service

Where Is the Knowbot Information Service?

- Knowbot Information Service servers currently run at nri.reston.va.us and sol.bucknell.edu
- Telnet to server host, port 185
- Mail to user netaddress@server-host

Documentation

- ? or help prints command summary
- man prints manual entry summary
- UNIX manual page
- IETF draft
- netaddress-users@sol.bucknell.edu (netaddress-users-request to join)

6.4 Privacy Enhanced Mail

Presentation by James Galvin/TIS

Privacy Enhanced Mail (PEM) is a mechanism whereby electronic mail can be augmented to provide the following security services on an end-to-end basis (originator to recipient):

- message integrity,
- message origin authentication and
- message confidentiality.

These services are provided by User Agents; no special processing requirements are imposed on the Message Transfer System. They can be provided on a site-by-site or user-by-user basis without impacting other services. Interoperability among heterogeneous electronic mail components is supported.

Trusted Information Systems, Inc., (TIS) in concert with BBN and RSA Data Security, Inc. (RSADSI), is developing an initial implementation of Privacy Enhanced Mail. It is being developed on Sun 3/60 and Sun SPARC workstations running SunOS 3 and SunOS 4, but should be compatible with other Berkeley derived UNIXs. The security services are being integrated into the Rand Message Handling (MH) user agent, although integration with other user agents is not precluded. The software will be made widely available to the Internet community this summer.

The PEM software is a realization of the specification defined by RFCs 1113, 1114 and 1115. The set of RFCs is the outgrowth of a series of meetings of the IRTF Privacy and Security Research Group (formerly the IAB Privacy Task Force). RFC 1113 defines the message encipherment and authentication procedures; RFC 1114 defines a certificate-based key management system; RFC 1115 specifies the algorithms, modes and identifiers used by the other RFCs. The RFCs are tightly-coupled with the DES and RSA encryption algorithms, although the use of other algorithms is not precluded.

TIS currently has a "closed" version of PEM operational now. This version is limited to use by a contained community supporting its own certification authority. The "open" version, to be released to the Internet community, will support a national certification authority and will allow unrelated originators and recipients to correspond freely. For further information readers may contact either Jim Galvin <galvin@tis.com> or Dave Balenson <balenson@tis.com>, both at (301) 854-6889.

Privacy Enhanced Mail: The TMail Implementation

James M. Galvin

Trusted Information Systems

May 3, 1990

- Introduction
- Privacy Enhanced Mail
- TMail Project
- Summary

- What is PEM? What is TMail? Who are the players?
- An introduction to PEM.
- Our implementation, status: work completed, in progress, to be done.
- Schedule and availability.

- Privacy and Security Research Group
- RSA Data Security, Inc.
- BBNCC

- Trusted Information Systems

David Balenson
Glenn Benson
Pam Cochrane
Sheila Haghighat

Privacy Enhanced Mail (PEM)

- PEM is specified by the following RFCs:

  1113: Message Encipherment and Authentication Procedures

  1114: Certificate Based Key Management

  1115: Algorithms, Modes and Identifiers

- PEM is a mechanism whereby electronic mail can be augmented to provide the following security services on an end-to-end basis:

  – message integrity

  – message origin authentication

  – message confidentiality

If assymmetric cryptosystems are used, non-repudiation is also supported.

End-to-end means originator to recipient user agent, independent of the message transfer system.

- Message Integrity - the property that data has not been altered or destroyed in an unauthorized manner.
  - Integrity is protected by the application of a Message Integrity Code algorithm to the message, the *signing* of the resulting value and the inclusion of the signed value in the PEM headers.

Signing is accomplished by the use of an asymmetric (public key) cryptographic algorithm. A user has a pair of keys that are inverses of each other. One, the public key, is made generally available, in order that others may use it to send encrypted mail to the user. The other, the private key is kept secret, in order to support message signing and the decryption of received messages.

- Message Origin Authentication - the corroboration that the originator of the message is as claimed.

  Since the message integrity code (MIC) is *signed* with the originator's private key, only the originator could have sent the message, if the MIC is verified by the recipient.

Integrity without authentication is not supported; the two services are tightly coupled.

- Message Confidentiality - the property that the data is not made available or disclosed to unauthorized individuals, entities or processes.

A data encrypting key (DES) is generated. The message is encrypted with this key. The key is included in the PEM headers, once for each recipient, with each encrypted with the public key of the respective recipient.

There is a Recipient-ID field preceding each key, so that each recipient can find its respective copy of the key.

There are four steps to realizing an outbound PEM message:

1. Local Form
2. Canonical Form
3. Authentication and Encipherment
4. Printable Encoding

Canonical form is analogous to the inter-SMTP representation is defined in RFC 821 and RFC 822. (ASCII with CR-LF delimiters)

- Services are provided by User Agents; no special processing requirements are imposed on the Message Transfer System.
- Interoperability among heterogenous mail components is supported.
- Public keys are distributed via certificates.

Services can be provided on a site-by-site or user-by-user basis without impacting other services.

Interoperability is guaranteed by the canonical form of the message.

A certificate contains a user's identity and the user's public key. It is signed by an issuer. The issuer's certificate, in turn, contains the issuer's public key but is signed by a well-known issuer.

TMail Project

- TMail is a realization of the PEM specification.
- Organized around 5 major components.
- Two versions: OPEN versus CLOSED.
- User agent and message transfer agent dependencies.

- The 5 major components.
  - Crypto Administrator Interface (CAI)
  - Local Key Manager (LKM)
  - Key Distribution Center (KDC)
  - Send Crypto Controller (SCC)
  - Receive Crypto Controller (RCC)

- Crypto Administrator Interface (CAI) is used by an organization to perform the following functions:
  - register users
  - delete registered users
  - other administrative functions

  Access control is enforced by the LKM.

- Local Key Manager (LKM)
  - Primarily a cache of often used certificates.
  - As a matter of policy, private keys may not leave the perimeter of the LKM. Therefore, the LKM provides encryption and decryption services using private keys.

- Local Key Manager (LKM), continued
  - Two forms of access control are supported:
    * users may operate on the themselves
    * crypto administrators may operate on anyone

    A form is administratively set. They are not mutually exclusive.

- Local Key Manager (LKM), continued
  - Dependent on a number of libraries:
    * certificate interface
    * cryptographic interface
    * data interface
    * distinguished name interface
    * key manager interface
    * certificate revocation list interface

- Local Key Manager (LKM), continued
  - Application library is provided for communicating with the LKM

- Key Distribution Center (KDC)
  - central repository of all certificates
  - issues all certificates
  - maintains list of expired or deleted certificates

- Send Crypto Controller (SCC)
  - Accepts an *ordinary* mail message.
  - Returns a PEM message.

  Available as a library so it is independent of a mail user agent.

- Receive Crypto Controller (RCC)
  - Accepts a PEM message.
  - Returns an *ordinary* mail message.

  Available as a library so it is independent of a mail user agent.

- Two versions: OPEN versus CLOSED.
  - OPEN version
    - open community with unrestricted size.
    - local key generation.
    - organizational notaries and short hierarchy of certification authorities.
    - ad hoc means for distribution of certificates and CRLs.

- Two versions: OPEN versus CLOSED, continued.
  - CLOSED version
    - closed, limited community.
    - local or centralized key generation.
    - centralized certification authority.
    - fully automated means for distribution of certificates and CRLs.
    - sensitivity labelling.

- User Agent and Message Transfer Agent (MTA) dependencies.

  We have integrated PEM with the Rand MH User Agent and are dependent on the Sendmail Message Transfer Agent. While MH may be independent of its MTA, the changes necessary to support PEM were only made to that portion of MH which uses Sendmail. We expect to integrate the SCC with MH such that this dependence will not be present in the OPEN version.

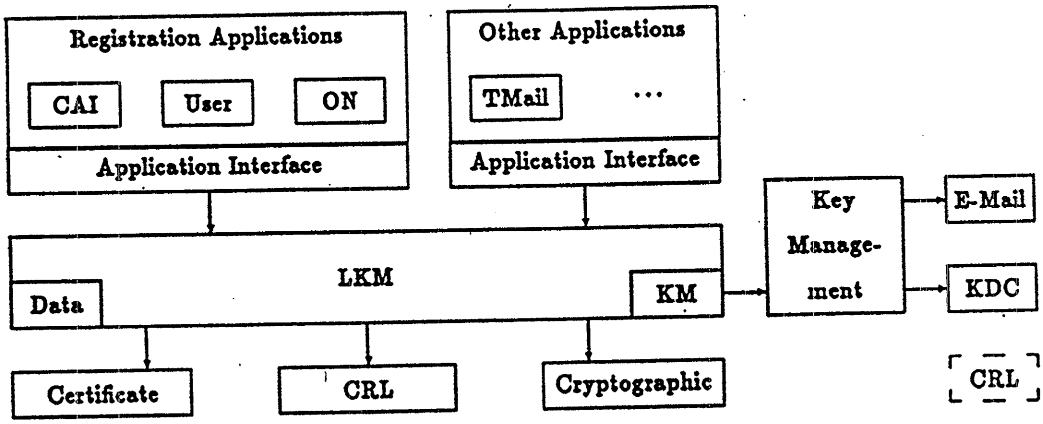

Registration Applications

CAI

User

ON

Application Interface

Other Applications

TMail

...

Application Interface

LKM

Data

KM

Key Manage-ment

E-Mail

KDC

CRL

Certificate

CRL

Cryptographic

Summary

- Beta testing to begin early in June.

- Wide availability late Summer.
  - SunOS3 and SunOS4
  - MC68020 and SPARC

- Performance statistics available at INTEROP90.

6.5 The FNC, NREN, and CCIRN

Presentation by Tony Villasenor

6.6 The FNC Engineering Planning Group

Presentation by Phill Gross/ NRI

Reported by Phill Gross

In late 1987, several U.S. agencies formed the Federal Research Internet Coordinating Committee (FRICC), under the auspices of the Federal Coordinating Committee for Science, Engineering and Technology (FCCSET). The goal was to coordinate agency network planning in laying the groundwork for a national networking infrastructure. To assist in designing and implementing the actual technical details, the FRICC formed the Federal Engineering Planning Group (FEPG).

The FRICC was an ad hoc group. In 1990, the FRICC was replaced by the more formally chartered Federal Network Council (FNC), reporting to the Office of Science and Technology Planning (OSTP). The FNC retains a very similar structure to the FRICC, with the exception that more agencies are involved, and separate functions are addressed in individual FNC working groups (Engineering and Operations, Security, and Research and Development). The FEPG is still the principal technical body of the FNC, but it now reports through the Engineering and Operations Working Group (EOWG), which sets FNC policy in the areas of network engineering and operations.

Tony Villasenor (Program Manager for the NASA Science Internet, Office of Space Science and Applications) presented an overview of the FNC, its objectives, its membership, and its working group structure. Mr. Villasenor chairs the FNC EOWG.

Mr. Villasenor's slides are quite thorough and self-explanatory. He includes a current history of the congressional NREN initiatives. He also describes how the FNC activities relate to the IETF and other international activities like the CCIRN.

I encourage anyone interested in the FNC or the NREN to review his excellent slides.

Phill Gross (CNRI) presented an overview of the FEPG and its current activities. The FEPG has been actively involved with domestic and international agency network interconnection, operational routing coordination in agency networks, OSI planning, and determining early agency requirements for NREN.

The FEPG has developed the notion of the Federal Inter-agency eXchange point (FIX) for interconnecting agency backbones (see slides). This is simply a central location in which various agency routers all connect to a single local area network. Currently, there are 2 such Fix's – FIX-East at the Suranet operations center in College Park Maryland, and FIX-West at Ames Research Center in California. There is consideration for installation of other FIX's, and for providing standard services at each FIX (e.g., NTP, DNS, etc). Coordination of inter-agency routing is generally done via the FIX's.

The FEPG has also helped plan two links to Europe for agency and general Internet traffic. The two planned links are from Goddard Space Flight Center to the University College of London, and from the ESnet in Princeton, New Jersey to the West German Research Network (WIN). This report gave only a brief overview the segregation of mission-oriented agency traffic from that meant for general Internet infrastructure. More detailed reports on these two interesting joint-agency links will be scheduled for a future IETF meeting after they become operational. The US-UK link is expected to be fully operational by September 1990. It is hoped the US-FRG link will be operational before the end of 1990.

A Status Report:

FNC & NREN & CCIRN

Presented to the

INTERNET ENGINEERING TASK FORCE
Pittsburgh Supercomputer Center
Pittsburgh, Pennsylvania
May 3, 1990

TONY VILLASENOR
Program Manager, NASA Science Internet
Office of Space Science and Applications
NASA

NASA Mission Organization

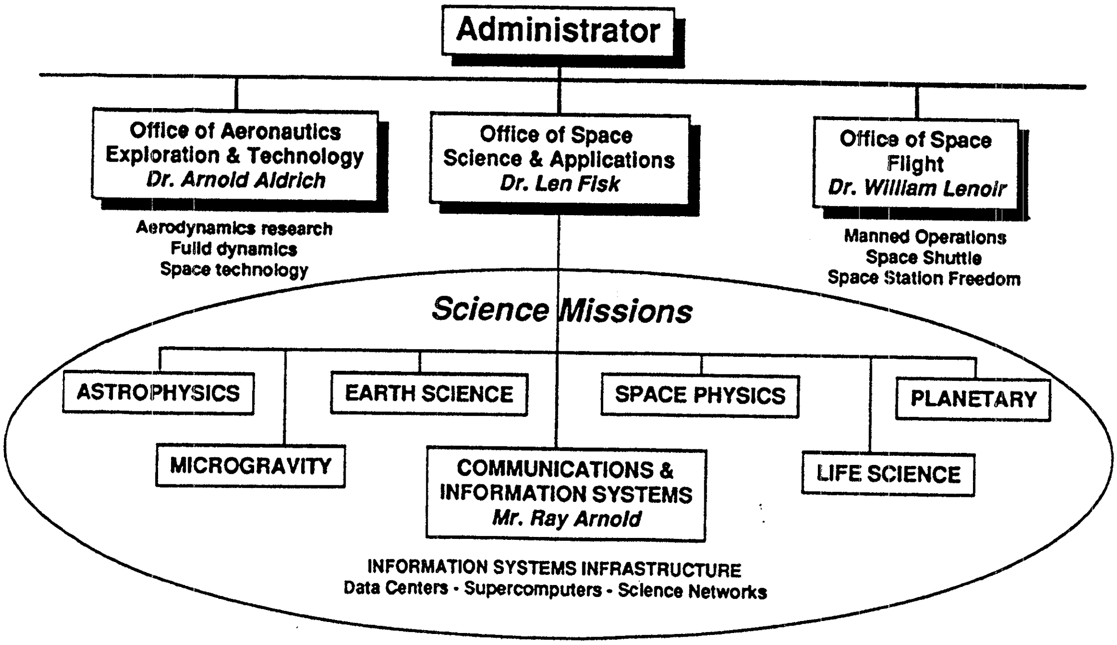

Administrator

Office of Aeronautics Exploration & Technology
Dr. Arnold Aldrich

Aerodynamics research
Fulid dynamics
Space technology

Office of Space Science & Applications
Dr. Len Fisk

Office of Space Flight
Dr. William Lenoir

Manned Operations
Space Shuttle
Space Station Freedom

Science Missions

ASTROPHYSICS
EARTH SCIENCE
SPACE PHYSICS
PLANETARY
MICROGRAVITY
COMMUNICATIONS & INFORMATION SYSTEMS
Mr. Ray Arnold
LIFE SCIENCE

INFORMATION SYSTEMS INFRASTRUCTURE
Data Centers - Supercomputers - Science Networks

FEDERAL NETWORKING COUNCIL

(FNC)

Federal Research Internet Coordinating Committee

(FRICC)

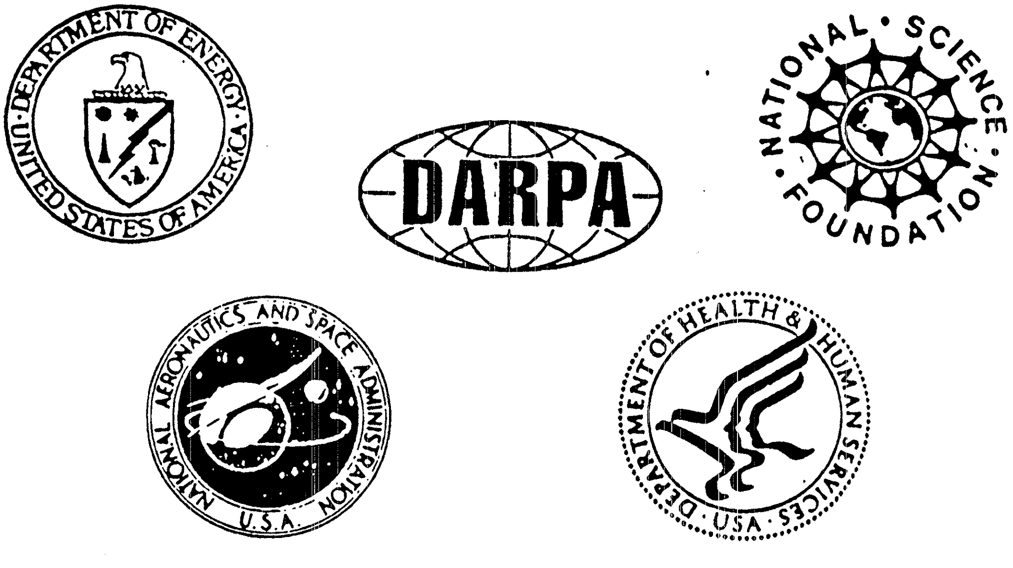
DEPARTMENT OF ENERGY · UNITED STATES OF AMERICA
DARPA
NATIONAL · SCIENCE · FOUNDATION
NATIONAL AERONAUTICS AND SPACE ADMINISTRATION · U.S.A.
DEPARTMENT OF HEALTH & HUMAN SERVICES · USA

FEDERAL NETWORKING COUNCIL

- FORMED BY FCCSET/Network SubCommittee Chair (JAN.4, 90)
- PURPOSE
  - PROVIDE [FEDERAL] POLICY DIRECTION FOR NREN VISION
  - COORDINATE ACTIVITIES & SERVICES OF FEDERAL NETS
  - ESTABLISH MECHANISMS TO ENSURE INTER-OPERATION
- I.E. - FNC WAS FORMED TO ESTABLISH AN INTERAGENCY FORUM AND LONG TERM STRATEGY TO OVERSEE THE OPERATION AND EVOLUTION OF THE INTERNET.
- MEMBERS FORMALLY DESIGNATED BY THEIR AGENCIES

IF NREN IS FUNDED, THE INTERNET WILL EVOLVE TO NREN

Federal Networking Council

Chairman - Dr. Charles Brownstein/NSF

GSA DARPA NIST DOD USGS NASA DCA NSF HHS NOAA DOE NTIA OSTP OMB

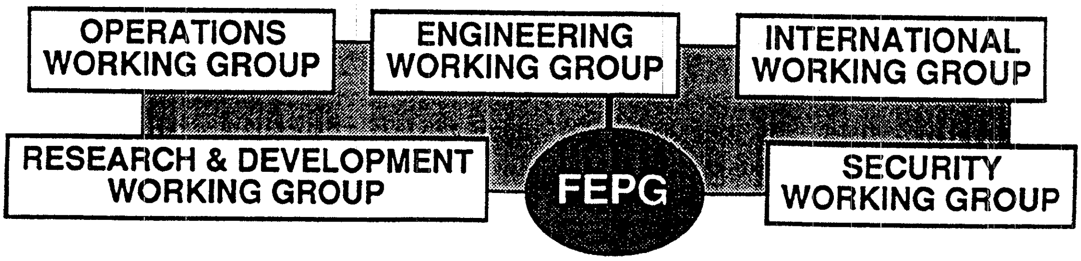
OPERATIONS WORKING GROUP
ENGINEERING WORKING GROUP
INTERNATIONAL WORKING GROUP
RESEARCH & DEVELOPMENT WORKING GROUP
FEPG
SECURITY WORKING GROUP

FNC Advisory Committee

FNC, coordinating with OSTP, will establish a charter and formal Advisory Committee representing industry and academia and the national user community; this Advisory Committee will work closely with the FNC to provide guidance in developing the NREN.

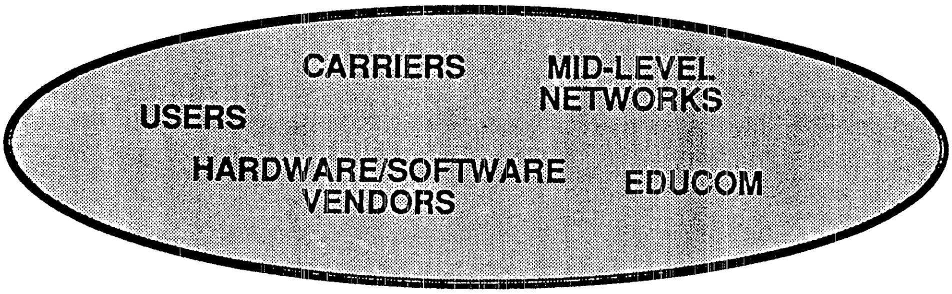

FNC & IETF

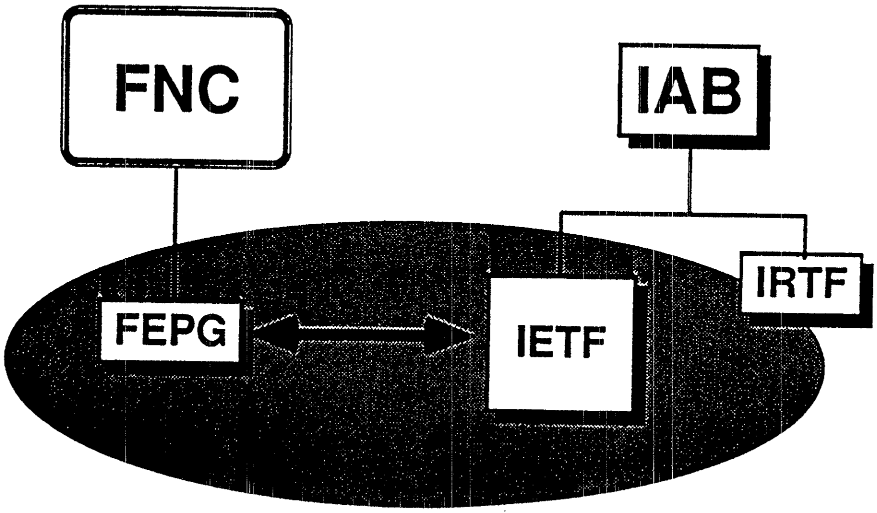

NATIONAL RESEARCH & EDUCATION NETWORK (NREN)

National Research & Education Network - NREN

HIGH PERFORMANCE COMPUTING INITIATIVE - HPCI

- Senator Gore (Bill S.1067) and Senator Johnston (Bill S.1976)
- Representative Walgren (Bill HR.3131)
- FRICC "NREN Program Plan"

NREN 5-YEAR 3-PHASE IMPLEMENTATION

Phase 1 - Interconnect Agency networks with 1.5 mbps backbone
Phase 2 - Upgrade multi-Agency backbone to 45 mbps
Phase 3 - Develop 3 gbps networking capability; research required

APPROACH

Build Internet to Interim NREN with NSF, DARPA, NASA, DOE, with placeholders for HHS, NOAA, USGS, NIST, etc.

Political Climate for NREN

May 3, 1990: afternoon

Commerce Committee unanimously passed S.1067 authorization and forwarded it to the full Senate for vote, probably in mid-May

Most "unusual": both Democratic and Republican sponsors

Support from *both* White House & Congress

Appears also to have growing support from education, libraries:

- not just for research and science
- for researchers in French literature, law, commerce, etc.

Congress wants to build a network that is compatible with industry; simple, so people can use it just like today's telephone services.

OMB: Good chances for FY92 program

1990 Chronology, up to May 3, 1990

Jan 1990	OSTP/Bromley called agency heads to support HPCI
	"HPCI looking good for FY 1992"
Jan-Feb 1990	FNC created => expanded & appointed FRICC
101st Congress	(3 separate Congressional authorization committees)
	Commerce: NSF & NASA - S.1067 (Senator Gore)
	Energy: DOE -S.1976 (Sen. Johnston, Gore, McClure)
	DARPA - part of Armed Services authorization bill
	Also, House version of Gore: HR.3131 (Rep. Walgren)
April 3, 1990	Commerce Authorization Committee approves S.1067; now forwarded to Senate for vote
April 4, 1990	House begins markup of Walgren bill - HR.3131
April 19, 1990	M.Nelson (Gore staffer) briefs FNC
May-July, 1990	Anticipate Senate markup & vote on S.1067 & S.1976

NREN ISSUES

WHITE HOUSE/OSTP/FCCSET

- Must insure total OSTP HPCI program; not just NREN
- Must have full agency agreement & [program] consistency

FNC

Different orientation between "Mission Agencies" and Infrastructure Agencies"

COMMERCIALICATION

- Need to define "commercialization"
- Define role of government in stimulating an NREN as well as associated technology transfer (gigabit protocols, switches, etc.)

TECHNOLOGY TRANSFER

Senate: no access restrictions in revised Gore version

House: wants to protect U.S. databases and resources

Result - HPCI to provide benefits to American companies, but not exclude other nations (viz, Free Trade Agreement with Canada)

NREN BUDGET

Summary of HPCI Funds Requested

	FY92	FY93	FY94	FY95	FY96
NREN					
Interagency Interim NREN	14	23	55	50	50
Gigabits R&D	16	27	40	55	60
COMPUTER SYSTEMS	55	91	141	179	216
ADVANCED SOFTWARE	51	90	137	172	212
BASIC RESEARCH	15	25	38	46	59

From: The Federal High Performance Computing Program, Executive Office of The President, 1989

Advanced Network Engineering Technology

PROTOCOL STANDARDIZATION
VERY HIGH SPEED SWITCHING
ADVANCED APPLICATIONS & SERVICES
POLICY-BASED INTERCONNECTION
NETWORK MANAGEMENT
SECURITY
ACCOUNTING

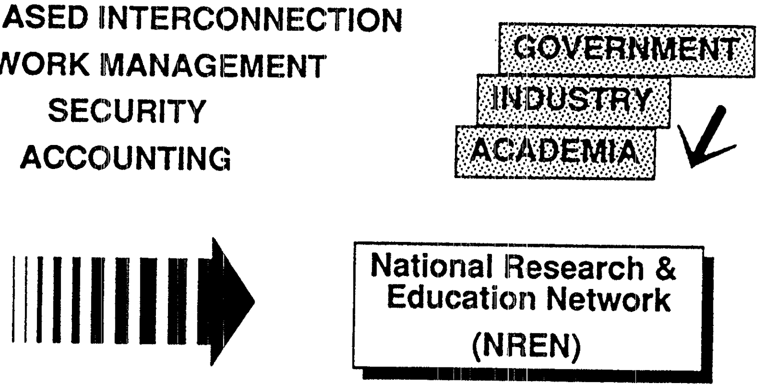

GOVERNMENT
INDUSTRY
ACADEMIA ✓

National Research & Education Network (NREN)

FY90 GIGABIT TESTBED PROGRAM (DARPA & NSF)

RESEARCH OBJECTIVES

- FOSTER GIGABIT APPLICATIONS
- EVALUATE ARCHITECTURAL ALTERNATIVES
- UNDERSTAND REQUIREMENTS FOR GIGABIT NETWORKS
- COLLABORATE WITH INDUSTRY, ESPECIALLY CARRIERS
- VEHICLE FOR APPLYING & TESTING RESEARCH RESULTS

STIMULATE GIGABIT APPLICATIONS

- DISTRIBUTED SUPERCOMPUTING
- LARGE SCALE MODELLING & SIMULATION
- INTERACTIVE VISUALIZATION

DARPA

Planned Gigabit Testbeds (4/90)

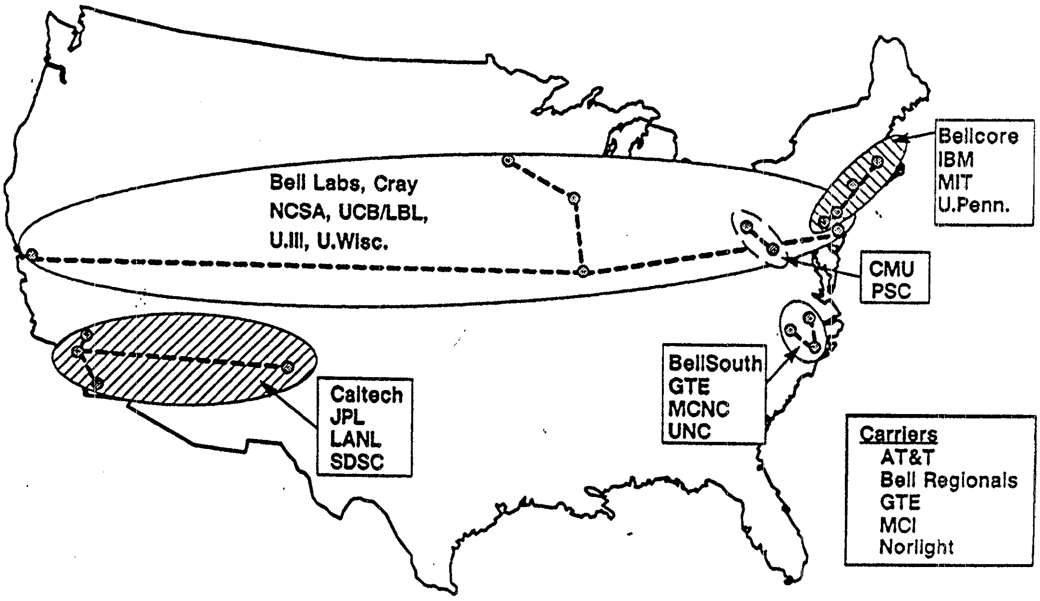

Bell Labs, Cray
NCSA, UCB/LBL,
U.Ill, U.Wisc.

Bellcore
IBM
MIT
U.Penn.

CMU
PSC

Caltech
JPL
LANL
SDSC

BellSouth
GTE
MCNC
UNC

Carriers
- AT&T
- Bell Regionals
- GTE
- MCI
- Norlight

4/90

INTERNATIONAL COLLABORATION

International Collaboration

NORTH AMERICAN COORDINATING COMMITTEE FOR INTERCONTINENTAL RESEARCH NETWORKING
Chair: Bill Bostwick

NACCIRN

NSF + DARPA + DOE + NASA + NIST + USGS + NOAA
(FNC agencies)
CREN (CSNET+BITNET) - FARNET
Canada (NRC, CANET, DND) & Mexico (CONACYT)

COORDINATING COMMITTEE FOR INTERCONTINENTAL RESEARCH NETWORKING
Chairs: US/Bill Bostwick and Europe/James Hutton

CCIRN

NACCIRN delegates (FNC agencies, IAB, Canada)
RARE + NORDUNET + JANET + DFN + EUNET + EARN + EURO-SPAN/HEPNET + ESA/ESTEC + RIPE
Observers (Japan, Australia)

CCIRN Issues

ACTIONS & ISSUES

- JANET gateway problems; role of RARE & U.S. to improve service.
- Role of U.S. Topology Engineering Working Group on international network architectures
- Approval of draft "CCIRN Policy on Intercontinental Leased Lines" from CCIRN/Canada meeting
- Seek CCIRN plan for the administration within Europe of IP network numbers, domain names, and autonomous system numbers.
- Define the current state of IP connectivity between Europe and North America. A TEWG task?
- Develop plan for improved U.S. connectivity to France, Scandinavia, Italy, and Greece.
- Begin discussions on international transit traffic issues.

NEXT MEETING • May 10-11 in Cannes, France

International [U.S.] Policy Issues

POLICIES FOR INTERCONNECTION: equitable sharing, contingent upon NREN and Pan-Europe multi-protocol networks (note: applies to infrastructure links only)

NOTE: FEPG policy on interconnection: e.g., international links should connect to an agency backbone, only one primary link between two countries, etc.

THEREFORE => *Interesting Topics for IETF, IAB, etc.*

- Monitoring and Accounting
- Routing & Topology Issues
- Security Protocols & Procedures *
- Interoperability
- Reliability & Risk
- Costing Strategies & Charging Algorithms

CHALLENGES & OPPORTUNITIES

The FUTURE

OBSERVATION: Multiplicity of developments!

- TCP/IP + OSI + Phase V + ANSI + GOSIP + + +
- FNC + U.S. "NREN"+ CEC/RARE, IXI + TISN + + +
- CLNP/CONP + ISDN + SONET + FDDI + + +

RESEARCH: More Opportunities than Ever!!

- Network Technology
- Network Management
- Network Applications & Services

NATIONAL INFRASTRUCTURE: Unprecedented changes!!!

- Roles of Government & Academia & Industry
- Commercialization & Privatization
- Increased use of Networking as a Researcy Tool
- Increased awareness of Networking as a Business
- New legislation required: domestic & international

Conclusion

YOU ARE TODAY
AT THE FOREFRONT
OF THE
EXPLODING TECHNOLOGY
IN
TELECOMMUNICATIONS

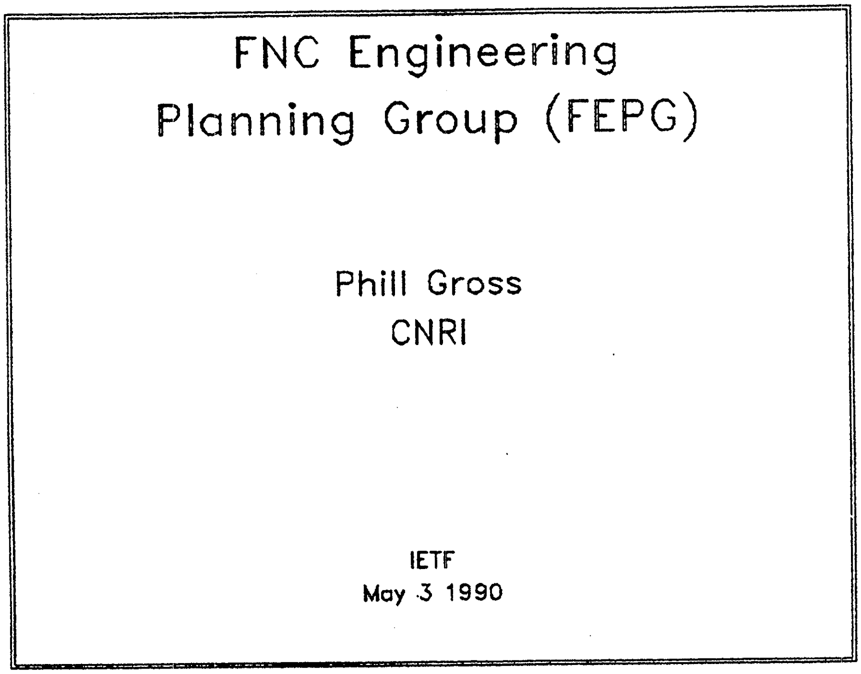

FNC Engineering Planning Group (FEPG)

Phill Gross
CNRI

IETF
May 3 1990

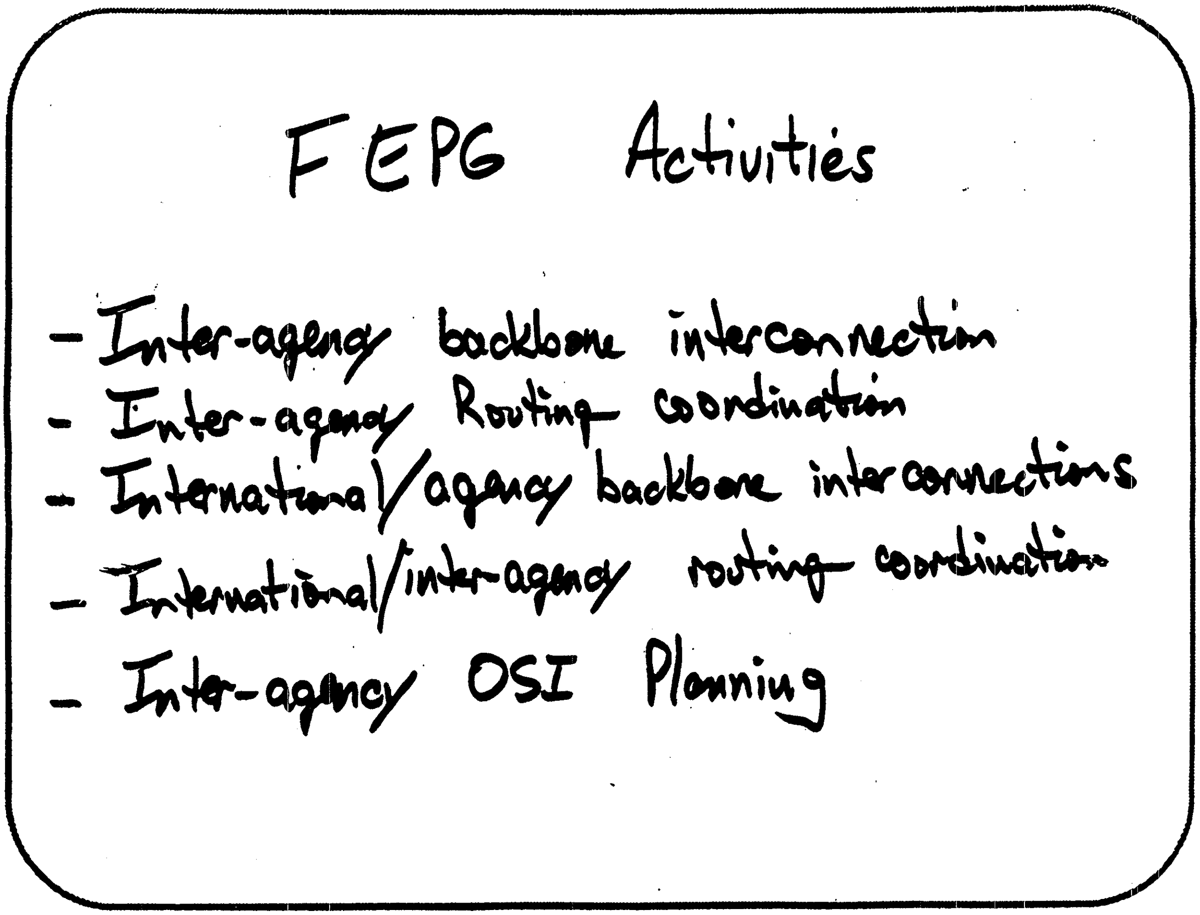

FEPG Activities

- Inter-agency backbone interconnection
- Inter-agency Routing coordination
- International/agency backbone interconnections
- International/inter-agency routing coordination
- Inter-agency OSI Planning

Coordinated Federal International Interconnectivity

An Example

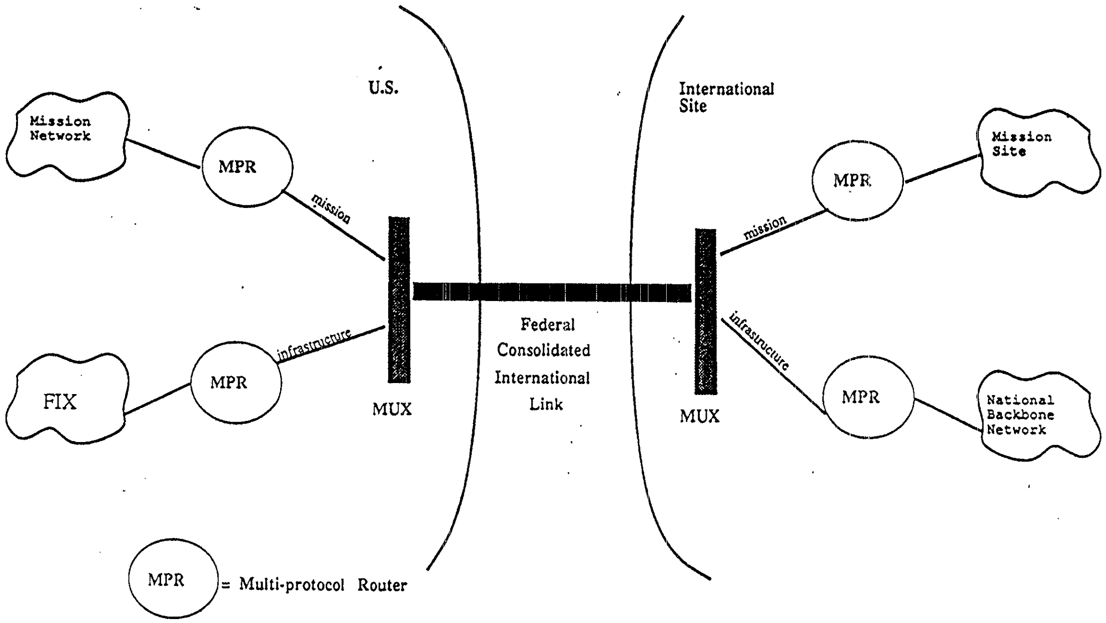

Federal Internetwork Exchange Point (FIX)

An Example

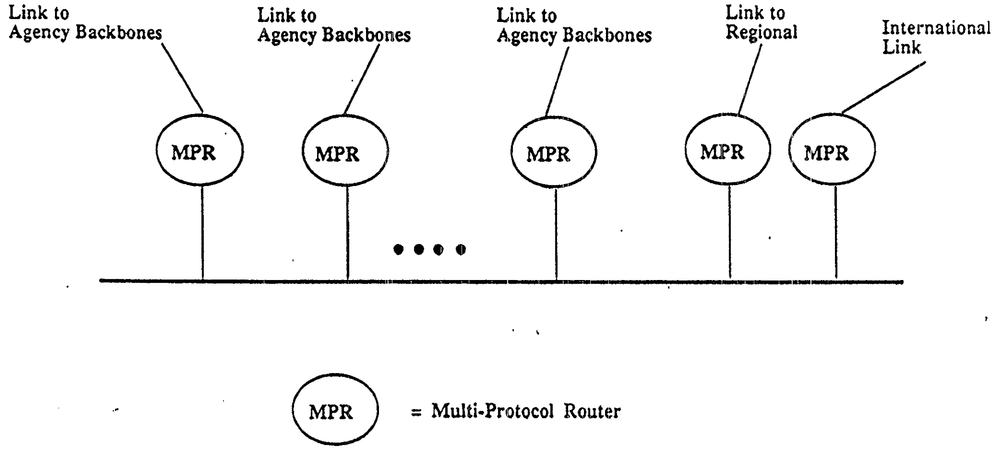

Federal Internetwork eXchange points (FIX)

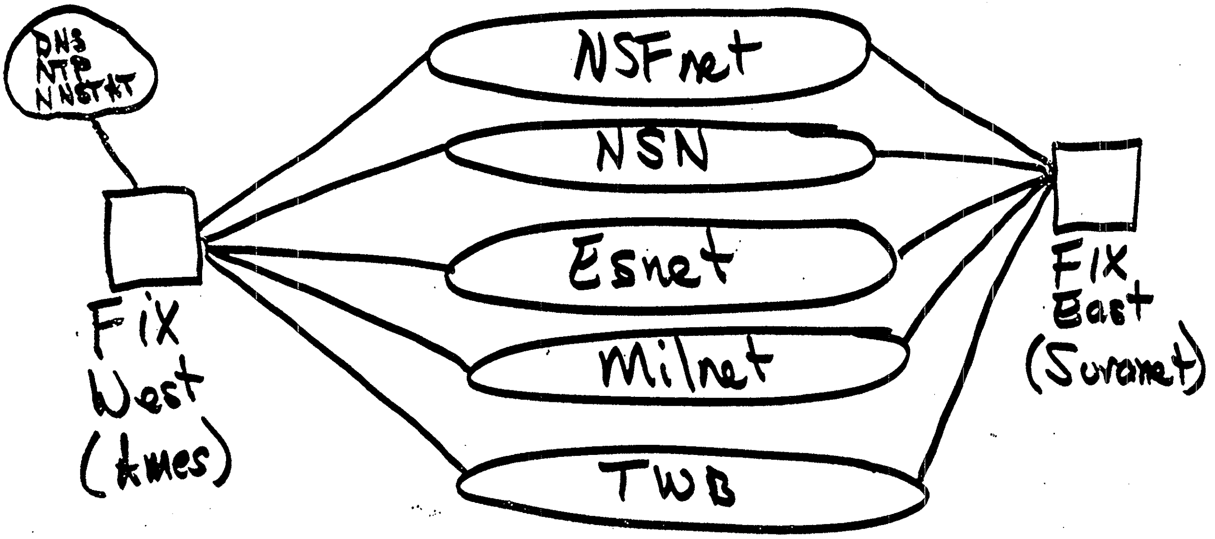

FIX West

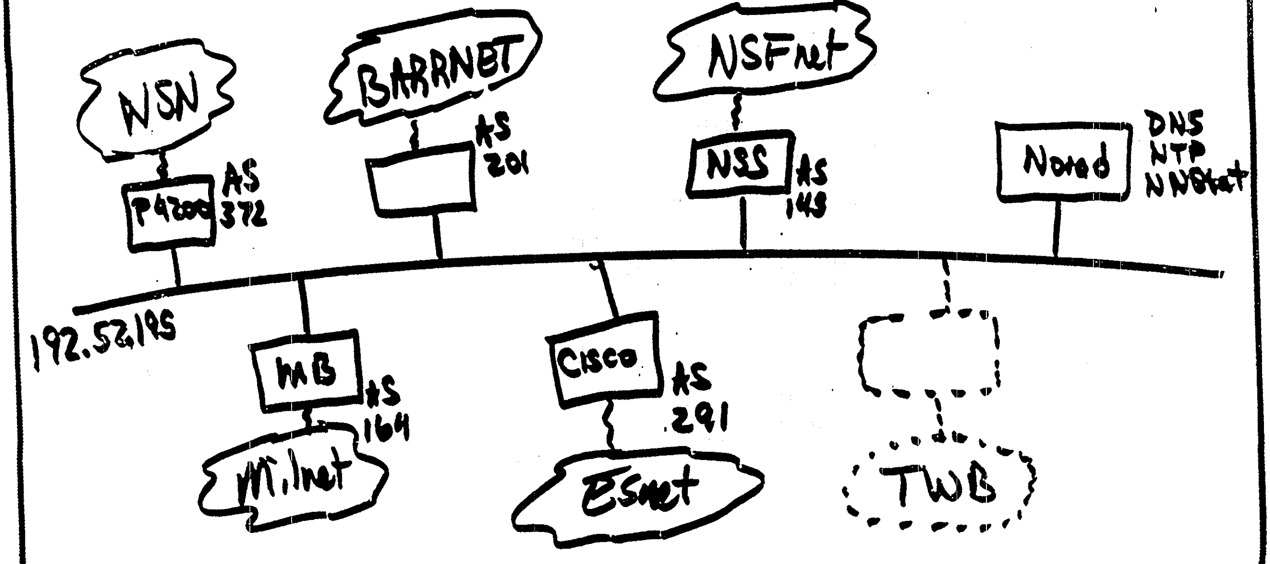

FIX East (Now)

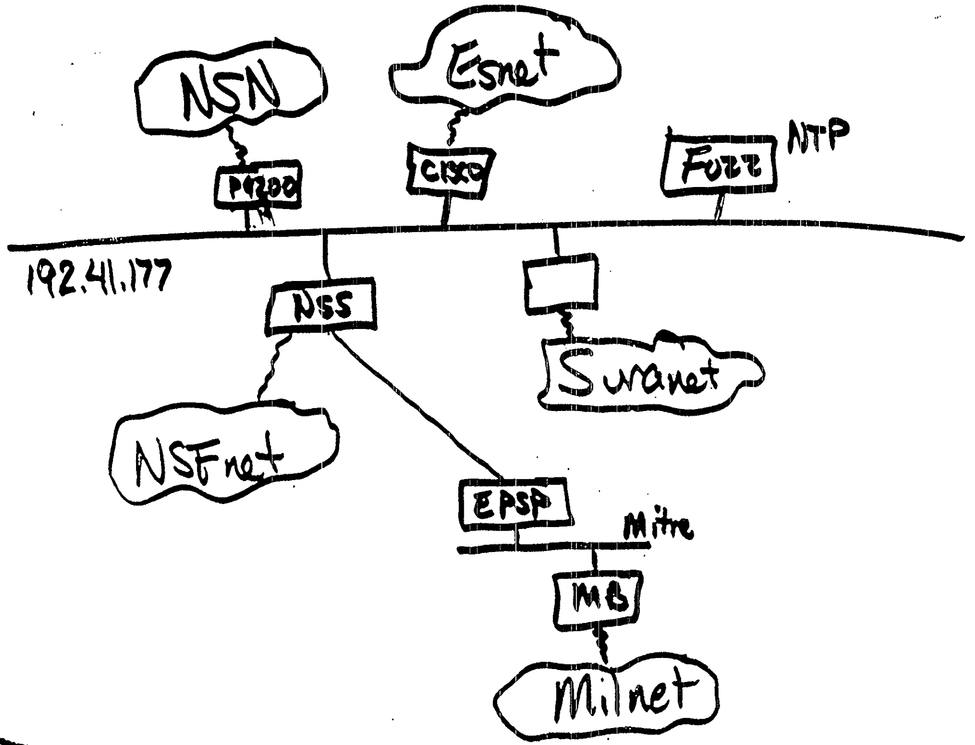

FIX East (RSN)

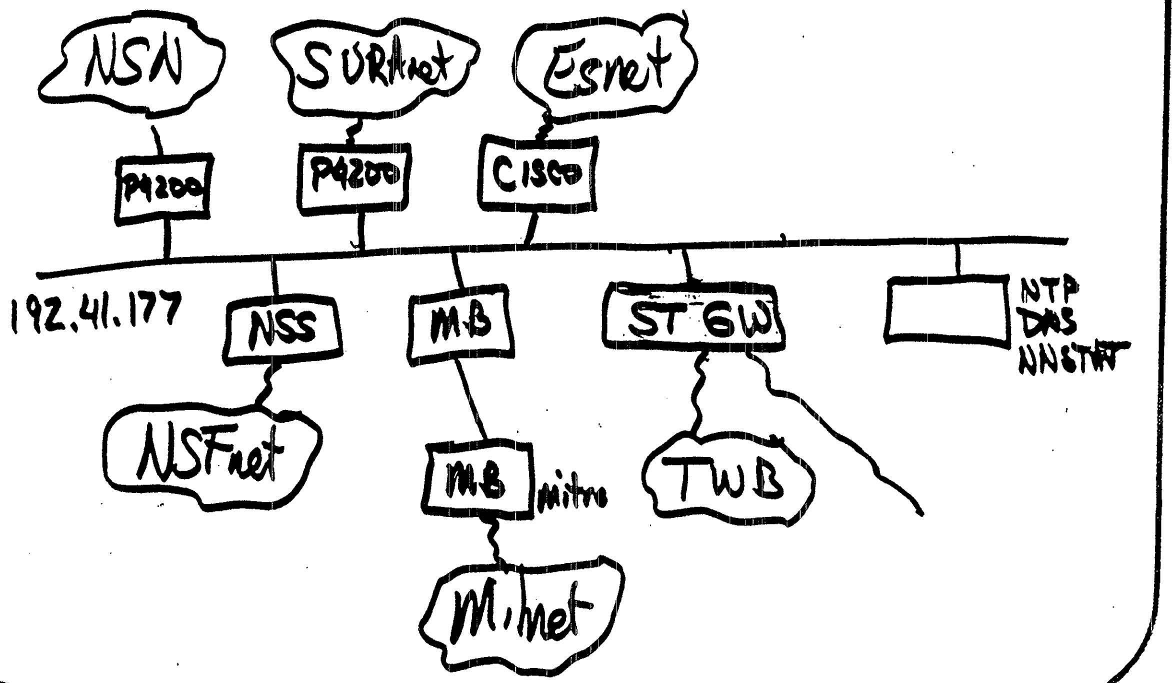

6.7 The Wide Area File System

Presentation by Philip Lehman/ Transarc

Wide area file service is a logical next addition to the set of services available on the Internet. A Wide Area File System includes transparent data sharing within and among potentially thousands of organizations, with tens of thousands of users per organization. A Wide Area File System utilizes a common space of file names and therefore a view of the file system that is consistent among users. This service has possibilities in as yet unexplored directions, including easy collaboration, personal mobility, information availability, and eventually a fundamental change to the work environment in-the-large. Such service must be truly easy to use, essentially hiding the existence of the underlying network from users. We describe an experiment that is about to be undertaken and that is based upon Transarc's AFS product (formerly the Andrew File System developed at Carnegie Mellon University).

Subscribers to such a Wide Area File System are, de facto, joining a federation of autonomous organizations (called "cells"), with common directory, protection, authorization, and transfer protocols and models. Requirements for such an endeavor include operation without central administration, networks supporting reliable and efficient long-distance communication, straightforward local management, high availability, and appropriate security mechanisms. The Wide Area File System consists of cells collaborating to provide a communal (very large) file name space, where the physical (machine) location of any given information is transparent to users.

The goal of the experiment is improved understanding of the viability of a shared data space, of opportunities for increased collaboration and of the scalability of these notions. The specific components of the project are a number of focus installations, a central cell (a clearinghouse and data repository, possibly called "GRAND.CENTRAL.ORG", not a central administrator) and research on both the technical and usage effects of a large distributed file system. Network effects to be studied include: effective throughput, data loss, latencies, availability, traffic, load, efficiency, and therefore the potential for further scaling. Additional technical measures include storage usage and economies, the degree of data sharing over small and large distances, and the use and reliability of authentication systems. Usage observations (via focus projects) include effects on collaboration and effectiveness, the issues of decentralized management, and the scaling of the logistics to tens or hundreds of sites.

The proposed project, therefore, hopes to provide a new level of network service, built upon new technologies, and a better understanding of the characteristics and uses of wide area networks.

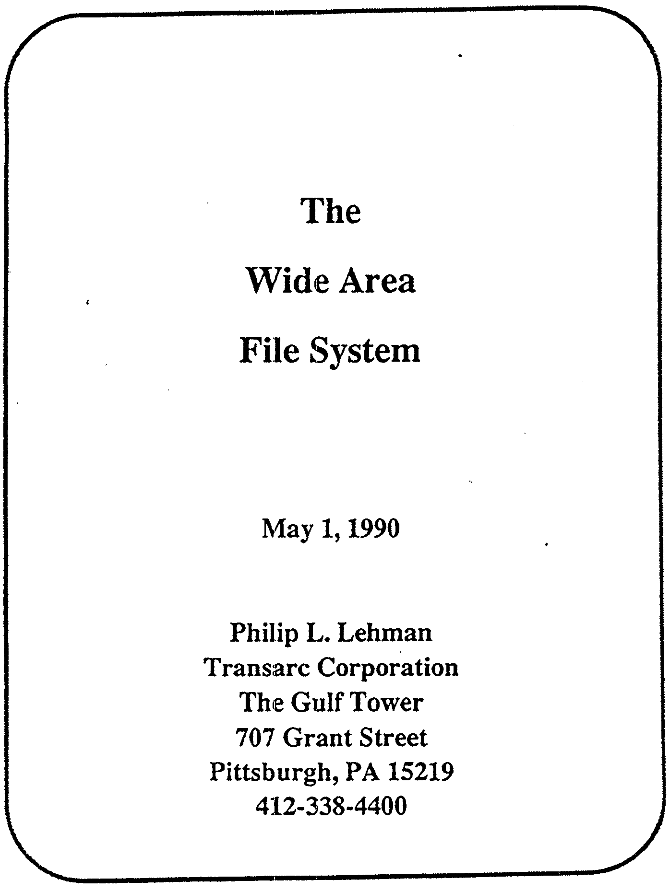

The Wide Area File System

May 1, 1990

Philip L. Lehman
Transarc Corporation
The Gulf Tower
707 Grant Street
Pittsburgh, PA 15219
412-338-4400

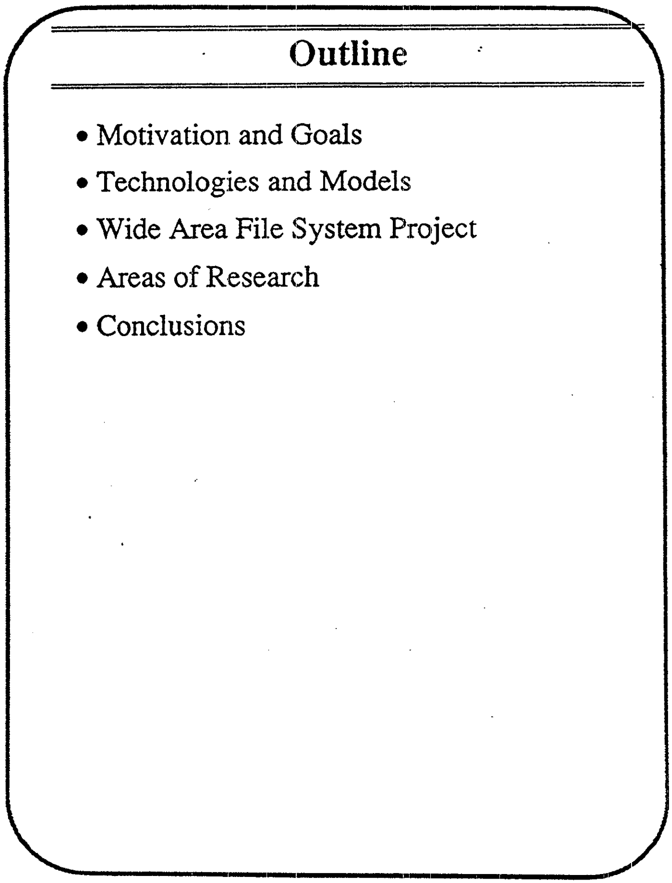

Outline

- Motivation and Goals
- Technologies and Models
- Wide Area File System Project
- Areas of Research
- Conclusions

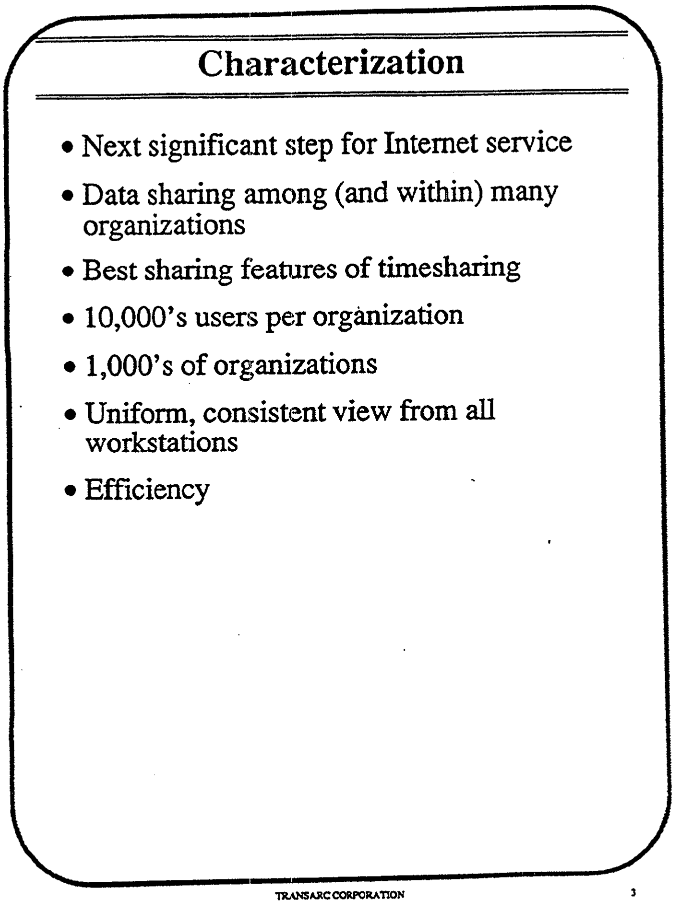

Characterization

- Next significant step for Internet service
- Data sharing among (and within) many organizations
- Best sharing features of timesharing
- 10,000's users per organization
- 1,000's of organizations
- Uniform, consistent view from all workstations
- Efficiency

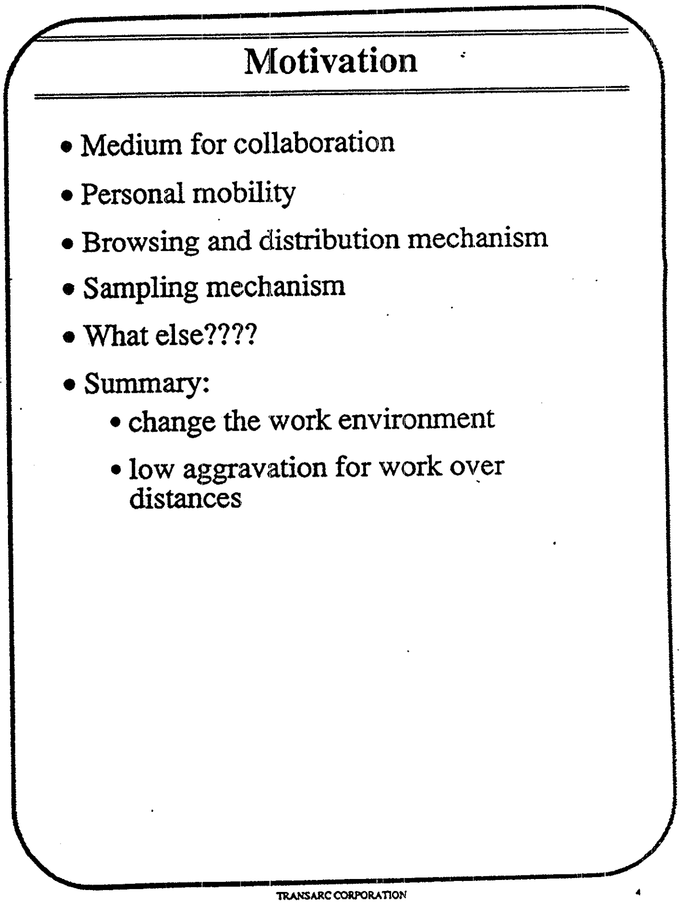

Motivation

- Medium for collaboration
- Personal mobility
- Browsing and distribution mechanism
- Sampling mechanism
- What else????
- Summary:
  - change the work environment
  - low aggravation for work over distances

History

1982	CMU Information Technology Center founded
1984	VICE-1: first campus deployment
1985	VICE-2: performance improvements (callbacks, threads)
1986	''Andrew''
1988	Nationwide File System Workshop
1989	AFS 3.0 beta: cell architectrure, long-haul, Kerberos, file chunks
1990	AFS 3.0 product: NFS interoperability, utilities (backup, monitoring)

Service Goal

Transparent access to authorized users

- Heterogeneous systems
- Multiple organizations
- Trans- and inter-continental distances
- True file system: hide existence of network
- Extension to existing network services: remote login, file transfer, mail

Practical Model

- Federations of autonomous organizations
- Common directory, protection, authorization, transfer protocols and models
- Bi-lateral agreement on access control
- Hence, no central administration required

Requirements

- Reliable, efficient communication
- Mechanisms to support easy use
- Not an administrative nightmare

Available Technologies

- Networking: T-1 and faster
- Protocols: RPC (streaming, authenticated)
- Security and authentication: Kerberos
- Distributed file systems: AFS

Technology Model

- Server Machines: disks, server processes
- Client Machines: logins, file use
- (One physical machine can be both.)
- Current implementation: Unix

Technology: Uniform Name Space

- Communal file space
- Federation of ''cells''
- Physical location transparency
- (Domain name used for cell name.)

Examples:

/afs/transarc.com/public/pll/talks/ietf0590.mss

/afs/cs.cmu.edu/usr/dn/proj/fs/data12

/afs/ir.stanford.edu/users/s/smith/tr/afs.tex

/afs/rice.edu/usr/almes/afs/notes.txt

Technology: Cells

- Administrative domains
- Conform to organizational considerations
- Autonomous (e.g. independent userid's)

Technology: Other

- Local caching of files on clients; efficient consistency mechanism
- Long-haul RPC techniques: streaming, error recovery
- Distributed file protocol
- Volumes for disk management
- Security: Kerberos, access control lists (ACL's)
- Replication for availability:
  - system data
  - file data

Wide Area File System Challenge

- Experiment with the viability of a shared data space
- Foster efficient cooperation by providing a new level of network service
- Will mechanisms scale?
- Theory: intercell effects happen only for intercell sharing

DARPA Project Components

- Installations
- Central Cell
- Research
  - Technical
  - Usage

Installations

- Growth in number
- 40 ''official'' participants
  - 25 ''server'' sites
  - 15 ''client-only'' sites
- Other additional participants
- Basis for much larger system
- NFS/AFS Gateway option for interoperation

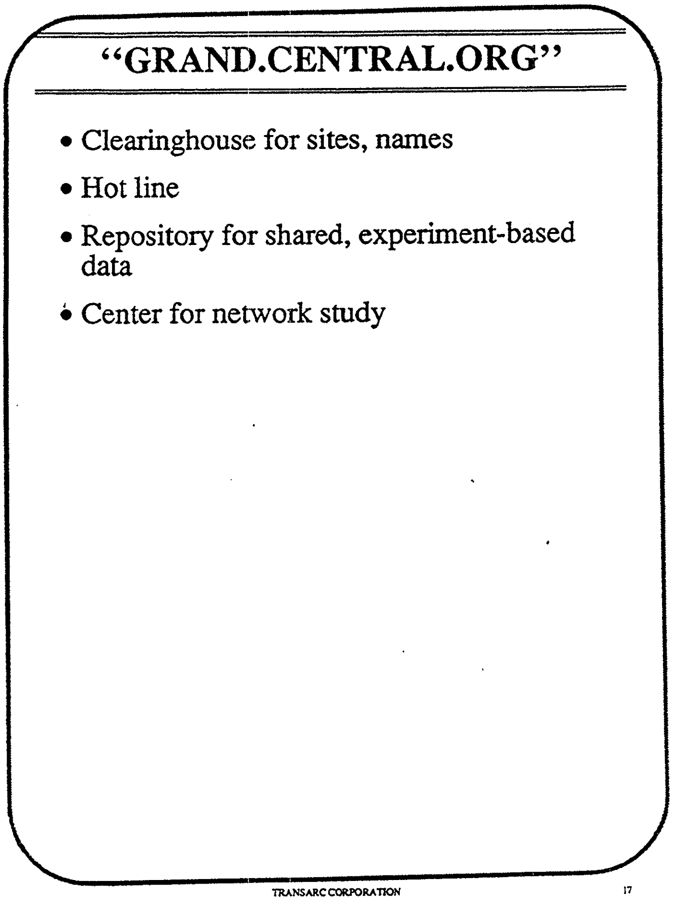

"GRAND.CENTRAL.ORG"

- Clearinghouse for sites, names
- Hot line
- Repository for shared, experiment-based data
- Center for network study

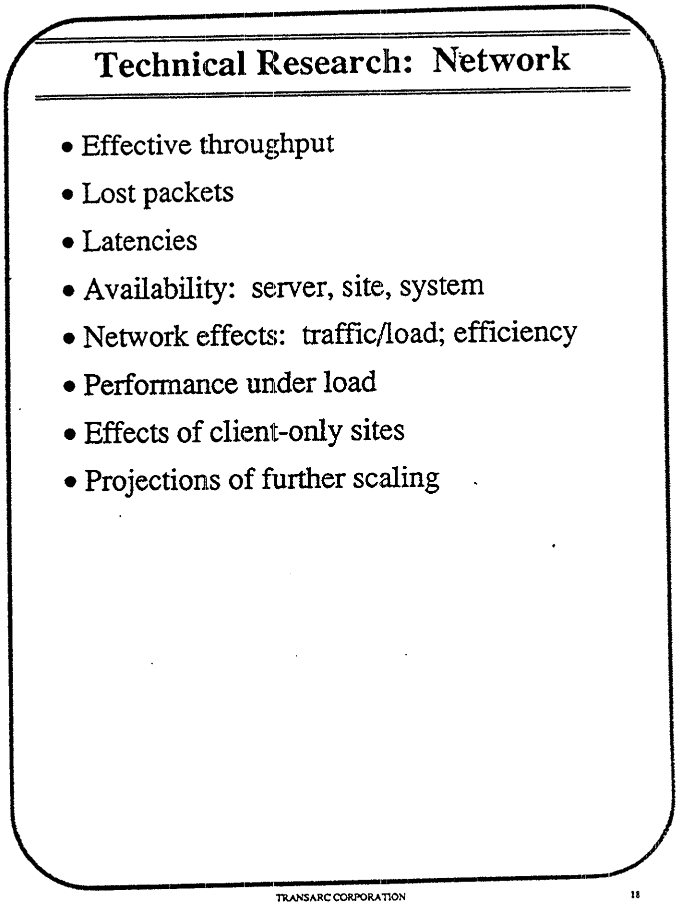

Technical Research: Network

- Effective throughput
- Lost packets
- Latencies
- Availability: server, site, system
- Network effects: traffic/load; efficiency
- Performance under load
- Effects of client-only sites
- Projections of further scaling

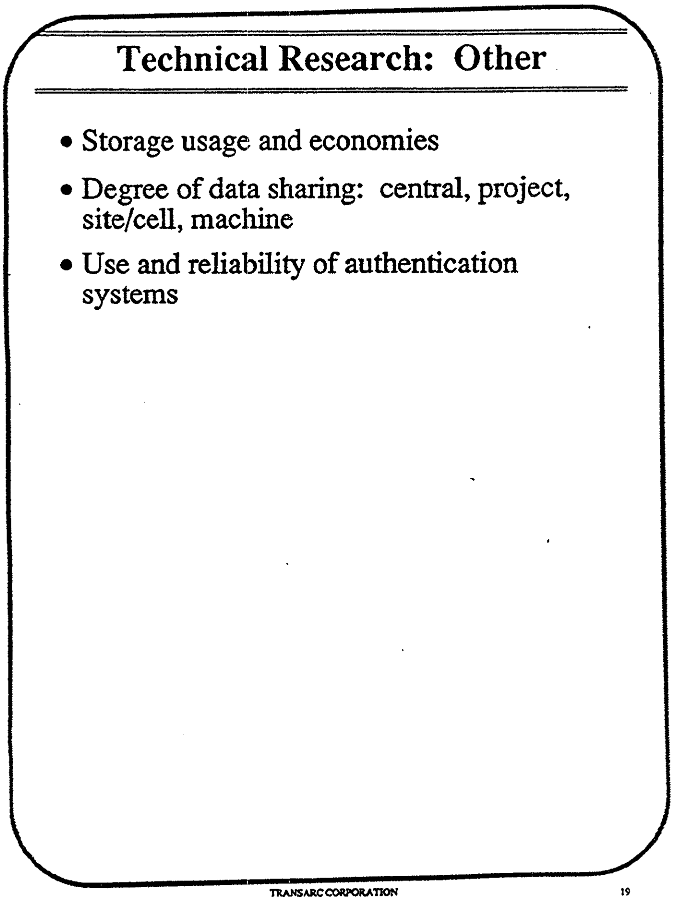

Technical Research: Other

- Storage usage and economies
- Degree of data sharing: central, project, site/cell, machine
- Use and reliability of authentication systems

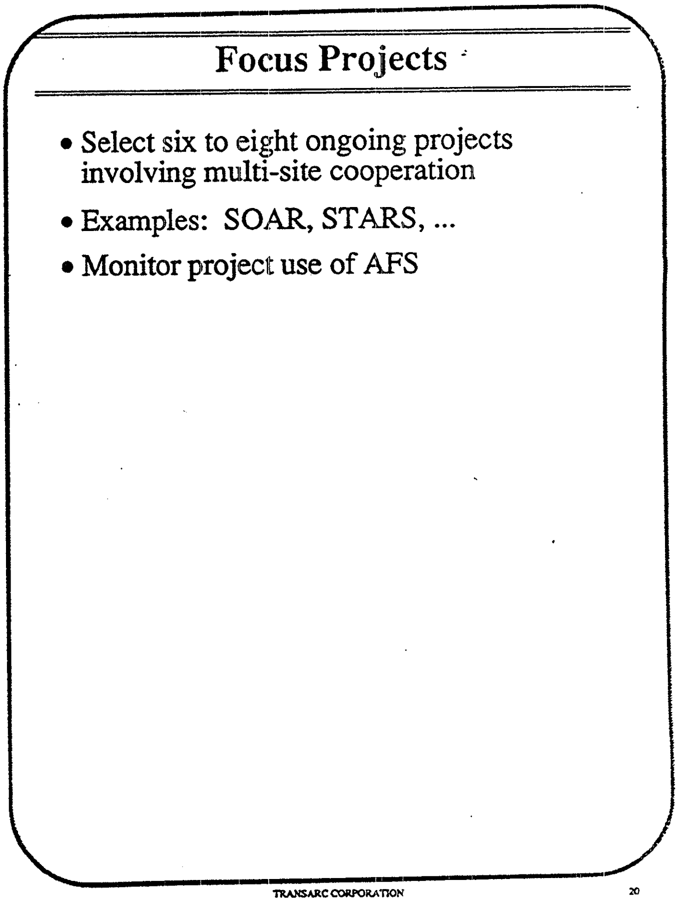

Focus Projects

- Select six to eight ongoing projects involving multi-site cooperation
- Examples: SOAR, STARS, ...
- Monitor project use of AFS

Usage Observations

- Interview/monitor participants: subjective/objective usage data
- How does collaborative work change?
- How do logistics scale to tens (hundreds) of sites?
- What is the function of the central site?
- Does decentralized management work?

Next Generation

- Proposed to O.S.F. as "DEcorum"
- RPC: NCS
- Kerberos V5
- POSIX semantics
- Protocol Translator for upward compatibility
- Per-file ACL's
- Other performance improvements
- RFC will be issued on distributed file protocol
- Intent:
  - Protocol refinement
  - Wide Area File System performance improvement
  - Creation of compatible utilities/enhancements

Calendar

- De facto: Wide Area File System exists today
- Project start: May, 1990
- End of Year 1: 24 project sites
- End of Year 2: 40 project sites
- Additional sites as well

Summary

- New level of service
- Built upon latest technologies
- Better understanding of characteristics and uses of wide-area networks

Summary

"No matter where you go... there you are."

Appendix A

Attendees

Jaap Akkerhuis
Mt. XINU
2560 Ninth Street
Berkeley CA, 94710
(W)415-644-0146
jaap@mtxinu.com

Hossein Alaee
3Com
2081 N. Shoreline Blvd.
Mountain View CA, 94043
(W)415-940-7648
hossein_alaee@3com.com

Guy Almes
Rice University
PO Box 1892
Houston TX, 77251-1892
(W)713-527-6038
almes@rice.edu

Philip Almquist
Consultant
214 Cole Street
San Francisco CA, 94117
(W)415-752-2427
almquist@jessica.stanford.edu

Stan Ames
Mitre
Burlington Road
Bedford MA, 01730
(W)617-271-3182
sra@mbunix.mitre.org

Douglas Bagnall
Hewlett-Packard
Apollo Division
Chelmsford MA, 01824
(W)508-256-6600
bagnall_d@apollo.hp.com

Thomas Bajzek
PREPnet
530 North Neville St.
Pittsburgh PA, 15213
(W)412-268-2637
twb@andrew.cmu.edu+

Fred Baker
Vitalink
6607 Kaiser Drive
Fremont CA, 94555
(W)415-794-1100
baker@vitalink.com

Patrick Barron
Transarc
Gulf Tower
Pittsburgh PA, 15219
(W)412-338-4451
pat_barron@transarc.com

Edward Berger
Pittsburgh Supercomputer Ctr.
4400 Fifth Ave.
Pittsburgh PA, 15213
(W)412-268-6481
eberger@b.psc.edu

Fred Bohle
Interlink Computer Sciences
10220 Old Columbia Road
Columbia MD, 21046
(W)301-290-8100
fab@saturn.acc.com

Dave Borman
Cray Research
1440 Northland Drive
Mendota Heights MN, 55120
(W)612-681-3398
dab@opus.cray.com

Leonard Bosack
cisco Systems
1525 O'Brien Drive
Menlo Park CA, 94025
(W)415-326-1941
bosack@mathom.cisco.com

Hans-Werner Braun
Merit Computer Network
University of Michigan
Ann Arbor MI, 48109
(W)313-763-4897
hwb@merit.edu

Terry Braun
Kinetics
1340 Treat Blvd
Walnut Creek CA, 94596
(W)415-947-0998
tab@kinetics.com

Pablo Brenner
Fibronics
Advanced Technology Ctr
Haifa , 31905 ISRA
(W)972 4 562645

Scott Brim
Cornell Theory Center
265 Olin Hall
Ithaca NY, 14853
(W)607-255-8686
swb@dainchi.tn.cornell.edu

Ronald Broersma
Naval Ocean Sys Ctr
Code 914
San Diego CA, 92152-5000
(W)619-553-2293
ron@manta.nosc.mil

Alison Brown
OARnet
Ohio Supercomputer Ctr
Columbus OH, 43212
(W)614-292-8100
alison@osc.edu

L. Allyson Brown
Department of Defense
C322
Ft. George G. Meade MD, 20755
(W)301-859-4502
allyson@umd5.umd.edu
labrown@dockmaster.ncsc.mil

Steven Bruniges
Retix
Alan Turing Road
Guildford SURREY, GU2 5YF ENGL
(W)44 483 300600

Theodore Brunner
Bellcore
445 South Street
Morristown NJ, 07960
(W)201-829-4678
`tob@thumper.bellcore.com`

Dave Burdelski
FTP Software
26 Princess Street
Wakefield MA, 01880-3004
(W)617-246-0900

Duane Butler
Network Systems
7600 Boone Avenue
Brooklyn Park MN, 55428
(W)612-424-4888
`dmb@network.com`

Glee Harrah Cady
Merit Computer Network
1075 Beal Avenue
Ann Arbor MI, 48109-2112
(W)313-936-3000
`ghc@merit.edu`

Ross Callon
Digital Equipment
550 King Street
Littleton MA, 01460-1289
(W)508-486-5009
`callon@bigfut.enet.dec.com`

C. Allan Cargille
U of Wisconsin
Computer Sciences Dept
Madison WI, 53706
(W)608-262-5084
`cargille@cs.wisc.edu`

Jeffrey Carpenter
U of Pittsburgh
Computing and Info Svcs
Pittsburgh PA, 15238
(W)412-624-6424
`j@unix.cis.pitt.edu`

Jeffrey Case
U of Tennessee
Dept of Computer Science
Knoxville TN, 37996
(W)615-974-0822
`case@utkux1.utk.edu`

Stephen Casner
USC
4676 Admiralty Way
Marina del Rey CA, 90292
(W)213-822-1511
`casner@venera.isi.edu`

Isidro Castineyra
BBN Communications
150 Cambridge Park Drive
Cambridge MA, 02140
(W)617-873-6233
isidro@bbn.com

John Cavanaugh
NCR Comten
Network Products Division
St. Paul MN, 55113
(W)612-638-2822
john.cavanaugh@stpaul.ncr.com

Vinton Cerf
Natl Research Initiatives
1895 Preston White Drive
Reston VA, 22091
(W)703-620-8990
cerf@isi.edu

Martina Chan
Motorola
1299 E Algonquin Road
Schaumburg IL, 60194
(W)708-576-1957
mchan@mot.com

A. Lyman Chapin
Data General
4400 Computer Drive
Westborough MA, 01580
(W)508-870-6056
lyman@merit.edu

Samir Chatterjee
Nynex
500 Westchester Avenue
White Plains NY, 10604
(W)914-683-2344
samir@nynexst.com

Andrew Cherenson
Silicon Graphics
2011 N. Shoreline Blvd.
Mountain View CA, 94039-7311
(W)415-962-3486
arc@sgi.com

J. Noel Chiappa
IESG
708 E. Woodland
Grafton VA, 23692
(W)804-898-7663
jnc@ptt.lcs.mit.edu

George Clapp
Ameritech Services
Gould Center, Building 40
Rolling Meadows IL, 60008
(W)708-806-8318
meritec!clapp@bellcore.bellcore.com

Danny Cohen
USC
4676 Admiralty Way #1000
Marina del Rey CA, 90292
(W)213-822-1511
cohen@isi.edu

Richard Colella
NIST
Bldg 225
Gaithersburg MD, 20899
(W)301-975-3627
colella@osi3.ncsl.nist.gov

Rob Coltun
U of Maryland
Computer Science Dept.
College Park MD, 20742-2411
(W)301-454-2946
rcoltun@trantor.umd.edu

Michael Contino
Penn State U
Office of Telecomm
University Park PA, 16802
(W)814-863-0859
mac@psuvm.psu.edu

John Cook
Chipcom
Southborough Office Park
Southborough MA, 01772
(W)508-460-8900
cook@chipcom.com

Curtis Cox
NARDAC
Code 303 Bldg 196
Washington DC, 20374-1435
(W)202-433-4026
zk0001@nhis.navy.mil

Bruce Crabill
U of Maryland
Computer Science Center
College Park MD, 20742
(W)301-454-2946
bruce@umdd.umd.edu

Caroline Cranfill
Bell South Services
675 W. Peachtree St.
Atlanta GA, 30375
(W)404-420-8432
rcc@bss.com
rcc@blsouth.com

Dave Crocker
Digital Equipment
Network Systems Lab
Palo Alto CA, 94301
(W)415-688-1320
dcrocker@nsl.dec.com

Steve Crocker
Trusted Information Sys
3060 Washington Road
Glenwood MD, 21738
(W)301-854-6889
crocker@tis.com

Steve Crumb
Motorola
3701 E Algonquin Road
Rolling Meadows IL, 60196
(W)708-576-9920
scrumb@genesis.corp.mot.com

Tom Cummings
PREPnet
NIC
Pittsburgh PA, 15213
(W)412-268-7874
tc1r@andrew.cmu.edu

James Davin
MIT
Computer Science Lab NE43-507
Cambridge MA, 02139
(W)617-253-6020
jrd@ptt.lcs.mit.edu

Edward DeHart
Carnegie Mellon U
CERT - SEI
Pittsburgh PA, 15213-3890
(W)412-268-6179
ecd@cert.sei.cmu.edu

Steve Deering
Stanford U
1017 Mallet Court
Menlo Park CA, 94025
(W)415-321-0224
deering@pescadero.stanford.edu

Peter DiCamillo
Brown U
PO Box 1885
Providence RI, 02912
(W)401-863-7582
cmsmaint@brownvm.brown.edu

Wilson Dillaway
U of Delaware
Academic Computing Support
Newark DE, 19716
(W)302-451-8447
dillaway@vax1.udel.edu

Greg Dobrich
Pittsburgh Supercomputer Ctr
4400 Fifth Avenue
Pittsburgh PA, 15213
(W)412-268-4960
dobrich@a.psc.edu

Tom Easterday
Ohio State U
1971 Neil Avenue
Columbus OH, 43210
(W)614-292-4027
tom@nisca.ircc.ohio-state.edu

Robert Enger
Contel
P O Box 10814
Chantilly VA, 22021-0814
(W)703-818-5555
enger@sccgate.scc.com

Hunaid Engineer
Cray Research
1400 Northland Drive
Mendota Heights MN, 55120
(W)612-681-3015
hunaid@opus.cray.com

Kent England
Boston U
Information Technology
Boston MA, 02215
(W)617-353-2780
kwe@buitb.bu.edu

Craig Everhart
Transarc
Gulf Tower
Pittsburgh PA, 15219
(W)412-338-4467
cfe@transarc.com

Roger Fajman
Natl Institutes of Health
Computer Center Bldg 12
Bethesda MD, 20892
(W)301-496-5181
raf@cu.nih.gov

Dino Farinacci
3Com
2081 N. Shoreline Blvd.
Mountain View CA, 94043
(W)415-940-7661
dino@buckeye.esd.3com.com

Dan Farmer
Carnegie Mellon U
CERT-SEI
Pittsburgh PA, 15213
(W)412-268-7090
df@sei.cmu.edu

Mark Fedor
Performance Systems Intl
Rensselaer Tech Park
Troy NY, 12180
(W)518-283-8860
fedor@psi.com

Dennis Ferguson
U of Toronto
5 King's College Road
Toronto ONTARIO, M5S 1A4 CANA
(W)416-978-2455
dennis@gw.ccie.utoronto.ca

Metin Feridun
BBN Systems & Technologies
10 Moulton Street
Cambridge MA, 02138
(W)617-873-1870
mferidun@bbn.com

Michael Fidler
Ohio State U
1971 Neil Avenue, Rm 406
Columbus OH, 43210-1210
(W)614-292-4843
ts0026@ohstvma.ircc.ohio-state.edu

Kathy Fithen
PREPnet
530 N. Neville St.
Pittsburgh PA, 15213
(W)412-268-7870
kf1b@andrew.cmu.edu+

Dave Forster
Digital Equipment
PO Box 121
Reading, Berkshire , RG2 OTU ENGL
(W)0734 868711
forster@marvin.enet.dec.com

Richard Fox
Hughes LAN
950 Linden Ave. #208
Sunnyvale CA, 94086
(W)415-966-7924
sytek!rfox@sun.com

Karen Frisa
Novell
1340 Treat Blvd
Walnut Creek CA, 94596
(W)415-947-0998
karen@kinetics.com

Stanley Froyd
Advanced Computer Comm
720 Santa Barbara Street
Santa Barbara CA, 93101
(W)805-963-9431
sfroyd@salt.acc.com

Vince Fuller
Stanford U
BARRNet Networking Systems
Stanford CA, 94305
(W)415-723-6860
vaf@valinor.stanford.edu

James Galvin
Trusted Information Sys
3060 Washington Road
Glenwood MD, 21738
(W)301-854-6889
galvin@tis.com

Ella Gardner
Mitre
7525 Colshire Drive
McLean VA, 22102-3481
(W)703-883-5826
epg@gateway.mitre.org

Fred Glover
Digital Equipment
110 Spitbrook Road
Nashua NH, 03062
(W)603-881-0388
fglover@decvax.dec.com

Martin Gross
DCA/DCEC
1860 Wichle Avenue
Reston VA, 22090-5500
(W)703-437-2165
gross@polaris.dca.mil

Phill Gross
Natl Research Initiatives
1895 Preston White Drive
Reston VA, 22091
(W)703-620-8990
pgross@nri.reston.va.us

Olafur Gundmundsson
U of Maryland
Dept. of Computer Science
College Park MD, 20742
(W)301-454-6497
ogud@cs.umd.edu

Robert Hagens
U of Wisconsin-Madison
Computer Science Dept.
Madison WI, 53706
(W)608-262-1017
hagens@cs.wisc.edu

Jack Hahn
SURAnet
U of Maryland
College Park MD, 20742
(W)301-454-5434
hahn@umd5.umd.edu

Tony Hain
Lawrence Livermore Natl Lab
PO Box 5509
Livermore CA, 94550
(W)415-422-4017
alh@eagle.es.net

Martyne Hallgren
Cornell Theory Center
265 Olin Hall
Ithaca NY, 14853-5210
(W)607-255-8686
martyne@tcgould.tn.cornell.edu

Brian Handspicker
Digital Equipment
550 King Street
Littleton MA, 01460
(W)508-486-7894
bd@vines.enet.dec.com

Richard Hart
Digital Equipment
110 Spit Brook Raod
Nashua NH, 03062-2698
(W)603-881-0418
hart@decvax.dec.com

Gene Hastings
Pittsburgh Supercomputer Ctr.
4400 Fifth Avenue
Pittsburgh PA, 15213
(W)412-268-4960
hastings@psc.edu

Robert Hinden
BBN Communications
50 Moulton Street
Cambridge MA, 02138
(W)617-873-3757
hinden@bbn.com

Don Hirsh
Meridian Technology
PO Box 2006
Ellisville MO, 63011
(W)314-394-1600
hirsh@magic.meridianpc.com

Russell Hobby
UC Davis
Computing Services
Davis CA, 95616
(W)916-752-0236
rdhobby@ucdavis.edu

Robert Hoffman
U of Pittsburgh
306 Alumni Hall
Pittsburgh PA, 15260
(W)412-624-8490
hoffman@cs.pitt.edu

Keith Hogan
Penril Data Communications
1300 Quince Orchard Blvd.
Gaithersburg MD, 20878-4106
(W)301-921-8600
uunet.uu.net.penril!keith

J. Paul Holbrook
Carnegie Mellon U
CERT-SEI
Pittsburgh PA, 15213-3890
(W)412-268-7720
ph@sei.cmu.edu

Greg Hollingsworth
Johns Hopkins U
Applied Physics Lab
Laurel MD, 20723
(W)301-953-6065
gregh@mailer.jhuapl.edu

Thomas Holodnik
Carnegie Mellon U
Networking & Comm, UCC-137
Pittsburgh PA, 15213
(W)412-268-2028
tjh@andrew.cmu.edu

Peter Honeyman
U of Michigan
CITI
Ann Arbor MI, 48103-4943
(W)313-763-4403
honey@citi.umich.edu

Jeffrey Honig
Cornell Theory Center
265 Olin Hall
Ithaca NY, 14853-5201
(W)607-255-8686
jch@tcgould.tn.cornell.edu

Michael Horowitz
Shiva
155 Second Street
Cambridge MA, 02141
(W)617-864-8500
mah@shiva.com

Kathleen Huber
BBN Communications
50 Moulton Street
Cambridge MA, 02138
(W)617-873-2520
khuber@bbn.com

Steven Hubert
U of Washington
Networks & Dist Computing
Seattle WA, 98195
(W)206-543-6384
hubert@cac.washington.edu

Steven Hunter
Lawrence Livermore Natl Lab
PO Box 808
Livermore CA, 94550
(W)415-423-2219
hunter@ccc.mfecc.arpa

Tim Hunter
Allegheny College
520 North Main Street
Meadville PA, 16335
(W)814-332-5307
thunter@allegum

Wendy Huntoon
Pittsburgh Supercomputer Ctr.
4400 Fifth Avenue
Pittsburgh PA, 15220
(W)412-268-4960
huntoon@a.psc.edu

Steve Huth
U of Pittsburgh
Computer Sciences & Information
Pittsburgh PA, 15260
(W)412-624-9369
skh@unix.cis.pitt.edu

Ole Jacobsen
Interop, Inc.
806 Coleman Avenue
Menlo Park CA, 94025
(W)415-325-9542
ole@csli.stanford.edu

Van Jacobson
Lawrence Berkeley Lab
One Cycloctron Road
Berkeley CA, 94720
(W)415-486-6411
van@helios.ee.lbl.gov

Brad Johnson
Open Software Foundation
11 Cambridge Center
Cambridge MA, 02142
(W)617-621-8849
bradcj@osf.org

Dan Jordt
U of Washington
170 Academic Computer Center
Seattle WA, 98105
(W)206-543-7352
danj@cac.washington.edu

Susie Karlson
Interop, Inc.
480 San Antonio Road
Mountain View CA, 94040
(W)415-941-3399
susie@milano.cisco.com

Phillip Karn
Bellcore
445 South Street
Morristown NJ, 07960
(W)201-829-4299
karn@thumper.bellcore.com

Frank Kastenholz
Racal Interlan
155 Swanson Road
Boxborough MA, 01719
(W)508-263-9929
kasten@europa.interlan.com

David Kaufman
Proteon
Two Technology Dr.
Westborough MA, 01581-5008
(W)508-898-2800
dek@proteon.com

Peter Kirstein
U College London
Department of Computer Science
London , N2 OAR ENGL
(W)44-1-380-7286
kirstein@cs.ucl.ac.uk

Stev Knowles
FTP Software
26 Princess Street
Wakefield MA, 01880-3004
(W)617-246-0900
stev@ftp.com

Alex Koifman
BBN Communications
50 Moulton Street
Cambridge MA, 02138
(W)617-873-8610
akoifman@bbn.com

Lee LaBarre
Mitre
Burlington Road
Bedford MA, 01730
(W)617-271-8507
cel@mbunix.mitre.org

Tracy LaQuey
U of Texas
Computation Center
Austin TX, 78712
(W)512-471-5046
tracy@emx.utexas.edu

Tony Lauck
Digital Equipment
550 King Street
Littleton MA, 01460-1289
(W)508-496-7644
lauck@tl.dec.com

Joseph Lawrence
Bellcore
331 Newman Springs Rd.
Redbank NJ, 07701
(W)201-758-4146
jcl@sabre.bellcore.com

Philip Lehman
Transarc
707 Grant Street
Pittsburgh PA, 15213
(W)412-338-4406
pll@transarc.com

James Leighton
Lawrence Livermore Natl Lab
PO Box 808
Livermore CA, 94551
(W)415-422-4025
jfl@nersc.gov

Mark Leon
NASA
Moffett Field
Mountain View CA, 94035
(W)415-604-6498
leon@nsipo.arc.nasa.gov

John Leong
Carnegie Mellon U
4910 Forbes Avenue
Pittsburgh PA, 15213
(W)412-268-6722
john.leong@andrew.cmu.edu

Joshua Littlefield
Cayman Systems
26 Landsdowne Street
Cambridge MA, 02139
(W)617-494-1999
josh@cayman.com

John LoVerso
Xylogics
53 Third Ave.
Burlington MA, 01803
(W)617-272-8140
loverso@xylogics.com

Kan Chei Loa
Motorola
2100 East Elliot Rd.
Tempe AZ, 85284
(W)602-897-5122
loa@sps.mot.com

Charles Lynn
BBN Systems & Technologies
10 Moulton Street
Cambridge MA, 02138
(W)617-873-3367
clynn@bbn.com

Yoni Malachi
Chipcom
PO Box 2014
Rehovot , 76120 ISRE
(W)972 8 460068
malachi@polya.stanford.edu

Louis Mamakos
U of Maryland
Computer Science Center - Syst
College Park MD, 20742
(W)301 454-2943
louie@trantor.umd.edu

George Marshall
Network Equipment Technologies
575 Chesapeake Drive
Redwood City CA, 94063
(W)415-366-9500
george@net.com

Marilyn Martin
CDNnet
CDNnet Headquarters, UBC
Vancouver BC, V6T 1W5 CANA
(W)604-228-6537
martin@cdnnet.ca

Tony Mason
Transarc
The Gulf Tower
Pittsburgh PA, 15219
(W)412-338-4400
mason@transarc.com

Matt Mathis
Pittsburgh Supercomputer Ctr
4400 Fifth Ave.
Pittsburgh PA, 15213
(W)412-268-3319
mathis@pele.psc.edu

Tony Mazraani
Washington U
Computer & Comm. Research Ctr.
St. Louis MO, 63130
(W)314-889-6106
tonym@flora.wustl.edu

Keith McCloghrie
Hughes LAN Systems
1225 Charleston Road
Mountain View CA, 94043
(W)415-966-7934
sytek!kzm@hplabs.hp.com

Leo McLaughlin
Wollongong Group
1129 San Antonio Road
Palo Alto CA, 94303
(W)415-962-7100
ljm@twg.com

Milo Medin
NASA Ames
Science Internet Project Office
Moffett Field CA, 94035
(W)415-490-9157
medin@nsipo.nasa.gov

Donald Merritt
Ballistic Research Lab
Attn: AMXBR-SECAD
Aberdeen Proving Ground MD, 21005-5066
(W)301-278-6808
don@brl.mil

David Miller
Mitre
Burlington Road
Bedford MA, 01730
(W)617-271-3993
dtm@ulana.mitre.org

Cyndi Mills
BBN Communications
150 CambridgePark Drive
Cambridge MA, 02140
(W)617-873-4143
cmills@bbn.com

Greg Minshall
Novell
1340 Treat Blvd.
Walnut Creek CA, 94596
(W)415-947-0998
minshall@optics.kinetics.com

Paul Mockapetris
USC
4676 Admiralty Way
Marina del Rey CA, 90292
(W)213-822-1511
pvm@isi.edu

Jeffrey Mogul
Digital Equipment
Western Research Labs
Palo Alto CA, 94301
(W)415-853-6643
mogul@decwrl.dec.com

Berlin Moore
PREPnet
530 N. Neville Street
Pittsburgh PA, 15213
(W)412-268-7873
bm24@andrew.cmu.edu+

Donald Morris
NCAR
PO Box 3000
Boulder CO, 80307
(W)303-497-1282
morris@ucar.edu

James Mostek
Cray Research
1440 Northland Drive
Mendota Heights MN, 55120
(W)612-681-3196
mostek@hall.cray.com

John Moy
Proteon
Two Technology Drive
Westborough MA, 01581-5008
(W)508-898-2800
jmoy@proteon.com

Mark Needleman
UC Berkeley
300 Lakeside Drive
Oakland CA, 94612-3550
(W)415-987-0530
mhnur@uccmvsa.bitnet

Oscar Newkerk
Digital Equipment
14475 NE 24th St.
Bellevue WA, 98007
(W)206-865-8913
newkerk@decwet.dec.com

Gerard Newman
San Diego Supercomputer Ctr
10100 John Jay Hopkins Drive
San Diego CA, 92186
(W)619-534-5076
gkn@sds.sdsc.edu

Torben Nielsen
U of Hawaii
Dept of ICS
Honolulu HI, 96822
(W)808-949-6395
torben@hawaii.edu

Matthew Nocifore
Drexel U
Office of Computing Services
Philadelphia PA, 19104
(W)215-895-2948
matthew@dupr.ocs.drexel.edu

Daniel Nydick
Pittsburgh Supercomputer Ctr.
Mellon Institute
Pittsburgh PA, 15213
(W)412-268-6886
nydick@psc.edu

John O'Hara
Digital/Project Athena
42 Chestnut St.
Andover MA, 01810
(W)617-253-6946
johara@athena.mit.edu

David O'Leary
U of Maryland
SURAnet
College Park MD, 20742
(W)301-454-8055
oleary@umd5.umd.edu

Lee Oattes
U of Toronto
Computer Services
Toronto ON, M5S 1C1 CANA
(W)416-978-5448
oattes@utcs.utoronto.ca

Zbigniew Opalka
BBN Communications
150 CambridgePark Drive
Cambridge MA, 02140
(W)617-873-2888
zopalka@bbn.com

Philippe Park
BBN Systems & Technologies
10 Moulton Street
Cambridge MA, 02238
(W)617-873-2892
ppark@bbn.com

Brad Parker
Cayman Systems
26 Landsdowne Street
Cambridge MA, 02139
(W)617-494-1999
brad@cayman.com

Paul Parker
Carnegie Mellon U
School of Computer Science
Pittsburgh PA, 15213
(W)412-268-7683
paul.parker@cs.cmu.edu

Michael Patton
MIT
545 Technology Square
Cambridge MA, 02139
(W)617-253-6061
map@lcs.mit.edu

Marc-Andre Pepin
CRIM
1550 De Maisonneuve West
Montreal QUEBEC, H3G 1N2 CANA
(W)514-848-3990
pepin@crim.ca

David Perkins
3Com
2081 N. Shoreline Blvd.
Mountain View CA, 94043
(W)415-694-2808
dave_perkins@3com.com

Drew Perkins
InterStream
824 Lilac Street
Pittsburgh PA, 15217
(W)412-422-9828
ddp@andrew.cmu.edu

Marsha Perrott
PREPnet
Carnegie Mellon U
Pittsburgh PA, 15213
(W)412-268-7871
mlp@andrew.cmu.edu+

Richard Pethia
Carnegie Mellon U
CERT-SEI
Pittsburgh PA, 15213-3890
(W)412-268-7739
rdp@sei.cmu.edu

Mike Petry
U of Maryland
Computer Science Center
College Park MD, 20742
(W)301-454-2943
petry@trantor.umd.edu

Tod Pike
Carnegie Mellon U
Software Engineering Institute
Pittsburgh PA, 15213-3890
(W)412-268-6814
tgp@sei.cmu.edu

David Piscitello
Bellcore
331 Newman Springs Road
Red Bank NJ, 07701
(W)201-758-2286
dave@sabre.bellcore.com

Paul Pomes
U of Illinois
1304 W Springfield Avenue
Urbana IL, 61801-2987
(W)217-333-6262
paul-pomes@uiuc.edu

Stephanie Price
CMC
125 Cremona
Santa Barbara CA, 93117
(W)805-968-4262
cmcvax!price@hub.ucsb.edu

Michael Reilly
Digital Equipment
Network Systems Lab
Palo Alto CA, 94301
(W)415-853-6593
reilly@nsl.dec.com

Yakov Rekhter
IBM
TJ Watson Research
Yorktown Heights NY, 10598
(W)914-945-3896
yakov@ibm.com

Joel Replogle
NCSA
152 Computing Applications Bldg.
Champaign IL, 61820
(W)217-333-1163
replogle@ncsa.uiuc.edu

Robert Reschly
US Army
Attn: SLCBR-SE-A (Reschley)
Aberdeen Proving Ground MD, 21005-5066
(W)301-278-6808
reschly@brl.mil

Joyce K. Reynolds
USC
4676 Admiralty Wy #1001
Marina del Rey CA, 90292-6695
(W)213-822-1511
jkrey@venera.isi.edu

Michael Roberts
EDUCOM
1112 16th Street NW
Washington DC, 20036
(W)202-872-4200
roberts@educom.edu

Ron Roberts
Stanford U
BARRNet
Stanford CA, 94305-4122
(W)415-723-3352
roberts@jessica.stanford.edu

Milt Roselinsky
CMC
125 Cremona Drive
Santa Barbara CA, 93117
(W)805-968-4CMC
cmcvax!milt@hub.ucsb.edu

Karen Roubicek
BBN Systems & Technologies
10 Moulton Street
Cambridge MA, 02138
(W)617-873-3361
roubicek@nnsc.nsf.net

Don Salvin
U of Pittsburgh
600 Epsilon Drive
Pittsburgh PA, 15238
(W)412-444-6324
dss@pitt.edu

Robert Sansom
Carnegie Mellon U
School of Computer Science
Pittsburgh PA, 15213
(W)412-268-5654
rds@cs.cmu.edu

Jonathan Saperia
Digital Equipment
550 King Street
Littleton MA, 01460-1289
(W)508-486-5542
saperia%tcpjon@decwrl.dec.com

Greg Satz
cisco Systems
1525 O'Brien Drive
Menlo Park CA, 94025
(W)415-326-1941
satz@cisco.com

Jeffrey Schiller
MIT
1 Amherst Street
Cambridge MA, 02139
(W)617-253-8400
jis@bitsy.mit.edu

Tim Seaver
Microelectronics Ctr of NC
PO Box 12889
Research Triangle Park NC, 27709
(W)919-248-1973
tas@mcnc.org

Mark Seger
Digital Equipment
550 King Street
Littleton MA, 01460
(W)508-486-5538
seger@mjs1.ogo.dec.com

Steve Senum
Network Systems
7600 Boone Avenue North
Minneapolis MN, 55428
(W)612-424-4888
sjs@network.com

Bob Shaw
PREPnet
530 N. Neville St.
Pittsburgh PA, 15213
(W)412-268-7870
rs6o@andrew.cmu.edu+

Jim Sheridan
IBM
166 East Shore Drive
Whitmore Lake MI, 48189
(W)313-393-6537
jsherida@ibm.com

Robert Shirey
Mitre
7525 Colshire Dr.
McLean VA, 22102-3481
(W)703-883-7210
shirey@mitre.org

Jim Showalter
DCA
1860 Wiehle Avenue
Reston VA, 22090-5500
(W)703-437-2580
gamma@mintaka.dca.mil

Robert Sidebotham
InterStream
370 Roup Ave.
Pittsburgh PA, 15232
(W)412-441-8342
bob@andrew.cmu.edu+

Dana Sitzler
Merit Computer Network
1075 Beal Avenue
Ann Arbor MI, 48109-2112
(W)313-936-3000
dds@merit.edu

Frank Slaughter
Shiva
155 Second Street
Cambridge MA, 02141
(W)617-864-8500
fgs@shiva.com

Mark Sleeper
Sparta
7926 Jones Branch Dr.
McLean VA, 22102
(W)703-448-0210

Pat Smith
Merit Computer Network
1075 Beal-NDSB
Ann Arbor MI, 48109
(W)800-66-Merit
psmith@merit.edu

Richard Smith
Datability Software Systems
322 Eighth Ave.
New York NY, 10001
(W)212-807-7800
smiddy@pluto.dss.com

Frank Solensky
Racal Interlan
155 Swanson Road
Boxborough MA, 01719
(W)508-263-9929
solensky@interlan.interlan.com

Ted Soo-Hoo
Data General
62 Alexander Dr.
Research Triangle Park NC, 27709
(W)919-549-8421
soo-hoo@dg-rtp.dg.com

Michael St Johns
US Dept of Defense
Attn: T41
Ft Meade MD, 20755
(W)301-688-6742
stjohns@umd5.umd.edu

Robert Stafford
Temple U
Computer/Math CIS Dept. , Rm. 300
Philadelphia PA, 19122
(W)215-787-6429
stafford@fac.cis.temple.edu

Mary K. Stahl
SRI International
Net Info Sys Ctr
Menlo Park CA, 94025
(W)415- 859-4775
stahl@nisc.sri.com

Martha Steenstrup
BBN Communications
150 CambridgePark Dr.
Cambridge MA, 02140
(W)617-873-3192
msteenst@bbn.com

Louis Steinberg
IBM
472 Wheelers Farms Rd
Milford CT, 06460
(W)203-783-7175
louiss@ibm.com

Tim Steiner
OARnet
State of Ohio
Columbus OH, 43215
(W)614-466-0747
steiner@oar.net

Steve Storch
BBN Systems & Technologies
10 Moulton St.
Cambridge MA, 02238
(W)617-873-3116
sstorch@bbn.com

Brad Strand
Cray Research
1440 Northland Dr.
Mendota Heights MN, 55120
(W)612-681-3249
bstrand@cray.com

Roxanne Streeter
NASA
Science Internet Project Office
Moffett Field CA, 94035
(W)415-694-4845
streeter@nsipo.arc.nasa.gov

Allen Sturtevant
Lawrence Livermore Natl Lab
PO Box 808
Livermore CA, 94550
(W)415-422-8266
sturtevant@ccc.nmfecc.gov

Zaw-Sing Su
SRI International
333 Ravenswood Ave.
Menlo Park CA, 94025
(W)415-859-4576
zsu@tsca.istc.sri.com

Thomas Swazuk
Temple U
Computer/Math, Room 763
Philadelphia PA, 19122
(W)215-787-6406
swazuk@fac.cis.temple.edu

Cal Thixton
NeXT
414 S. Craig
Dallas TX, 75380
(W)214-645-6398
cthixton@next.com

Ian Thomas
Chipcom
118 Turnpike Road
Southboro MA, 01772
(W)508-460-8900
ian@chipcom.com

Claudio Topolcic
BBN Systems & Technologies
10 Moulton Street
Cambridge MA, 02138
(W)617-873-3874
topolcic@bbn.com

Paul Tsuchiya
Bellcore
435 South Street
Morristown NJ, 07960
(W)201-829-4484
tsuchiya@thumper.bellcore.com

Kannan Varadhan
OARnet
Ohio Supercomputer Center
Columbus OH, 43212
(W)614-292-4137
kannan@oar.net

Gregory Vaudreuil
Corp for Natl Research Initiatives
1895 Preston White Drive, Suite 100
Reston VA, 22091 USA
(W)703-620-8990
gvaudre@nri.reston.va.us

John Veizades
Apple Computer
20525 Mariani Avenue
Cupertino CA, 95014
(W)408-974-2672
veizades@apple.com

Sudhanshu Verma
Hewlett-Packard
19420 Homestead Road
Cupertino CA, 95014
(W)408-447-3417
verma@hpindbu.cup.hp.com

Edward Vielmetti
U of Michigan
Dept. of Mathematics
Ann Arbor MI, 48109
(W)313-764-9497
emv@math.lsa.umich.edu

Anthony Villasenor
NASA & FNC & CCIRN
NASA Headquarters
Washington DC, 20546
(W)202-453-1495
Villasenor@nasa.gov

Peter Vinsel
Farallon Computing
2000 Powell St.
Emeryville CA, 94608
(W)415-596-9213
farcomp!pcv@apple.com

John Vollbrecht
U of Michigan
MERIT
Ann Arbor MI, 48109-2099
(W)313-763-1206
jrv@merit.edu

David Waitzman
BBN Systems and Technologies
10 Moulton Street
Cambridge MA, 02238
(W)617-873-4323
djw@bbn.com

Steve Waldbusser
Carnegie Mellon U
4910 Forbes Avenue
Pittsburgh PA, 15213
(W)412-268-6628
sw01@andrew.cmu.edu+

Y C Wang
NAT
21040 Homestead Road
Cupertino CA, 95014
(W)408-733-4530

Carol Ward
Westnet
U of Colorado
Boulder CO, 80309-0455
(W)303-492-5860
cward@naiad.colorado.edu

Bill Warner
State of Ohio
65 E. State Street
Columbus OH, 43215
(W)614-466-6683
warner@mps.ohio-state.edu

Jonathan Wenocur
Shiva
155 Second Street
Cambridge MA, 02141
(W)617-864-8500
jhw@shiva.com

Steve Willis
Wellfleet Communications
12 DeAngelo Drive
Bedford MA, 01730
(W)617-275-2400
swillis@wellfleet.com

Walter Wimer
Carnegie Mellon U
Networking & Communications
Pittsburgh PA, 15213-3890
(W)412-268-6252
ww0n@andrew.cmu.edu+

Linda Winkler
Argonne National Laboratory
Building 221, B-251
Argonne IL, 60439
(W)708-972-7236
b32357@anlvm.ctd.anl.gov

Dan Wintringham
OARnet
Ohio Supercomputer Center
Columbus OH, 43212
(W)614-292-0901
danw@igloo.osc.edu

John Wobus
Syracuse U
Computing and Network Services
Syracuse NY, 13244
(W)315-443-4324
jmwobus@suvm.acs.syr.edu

C. Philip Wood
Los Alamos Natl Lab
Network Engineering
Los Alamos NM, 87545
(W)505-667-2598
cpw@lanl.gov

David Wood
Mitre
7525 Colshire Drive
McLean VA, 22102
(W)703-883-6394
wood@gateway.mitre.org

Robert Woodburn
SAIC
CSEIC
Vienna VA, 22182
(W)703-734-9000
woody@saic.com

Richard Woundy
IBM
48 Wellington Road
Milford CT, 06460
(W)203-783-4308
rwoundy@ibm.com

Sze-Ying Wuu
JvNCnet
Consortium of Scientific Computing
Princeton NJ, 08543
(W)609-520-2000
wuu@nisc.jvnc.net

Mary Youssef
IBM
472 Wheelers Farms Rd.
Milford CT, 06460
(W)203-783-4338
mary@ibm.com

Aileen Yuan
Mitre
7525 Colshire Dr.
McLean VA, 22102
(W)703-883-7023
aileen@gateway.mitre.org

Chin Yuan
Pacific Bell
2600 Camino Ramon
San Ramon CA, 94583
(W)415-867-5016
cyuan@srv.pacbell.com