Proceedings of the Twentieth Internet Engineering Task Force

Washington University
March 11-15, 1991

Corporation
for
National
Research
Initiatives

ACKNOWLEDGEMENTS

The Twentieth Internet Engineering Task Force held in St. Louis, Missouri, was not only the largest to date, with 350 attendees, but also one of the most successful. In large part that success was do to the hard work and enthusiasm of Guru Parulkar, Martin Dubetz, the staff of the Office of the Network Coordinator and a host of student volunteers. We are very appreciative of their efforts on behalf of the IETF. We wish to recognize as well, the following for their generous contributions:

Southwestern Bell Telephone	T1 Connection and SMDS Connectivity
SBC Technology Resources, Inc.	T1 Connection and SMDS Connectivity
Jonathan Turner	Tour of Applied Research Laboratory
James Blaine	Tour of Electronic Radiology Laboratory
Apple Computer	Macintosh IIfx, IIsi, LaserWriter
cisco Systems	MGM Terminal Server
Digital Equipment Corporation	DECstation 5000, 3100, 4VT1200 VT100-compatible
Sun Microsystems	SPARCstation 2, SPARCstation 1, SPARCstation IPC

In addition there was an extensive list of individuals whose hard work contributed to the high quality of the Technical Presentation Sessions. Their efforts also deserve acknowledgement:

Deepinder Sidhu, Bill Johnston, Peter Ford, Bob Hinden, Philip Almquist, Bill Bostwick, Jonathan Turner, Rick Wilder, John Moy, Jeff Burgan, Rob Coltun, Dennis Ferguson, Sue Hares, Yakov Rekhter, Louis Steinberg, Kirk Williams and Andy Malis.

Many thanks also to Candice Moshos for her patient, creative and invaluable assistance in planning the logistics for the IETF.

Megan Davies/CNRI
Greg Vaudreuil/CNRI

Contents

Chairman's Message

St. Louis IETF Meeting

Once again, I must search for a new set of superlatives to describe our latest IETF hosts. I want to express my thanks to Washington University, especially Guru Parulkar, Martin Dubetz, and Kathy Atnip for the outstanding support and services they provided at the March 1991 IETF meeting. Thanks to their fine efforts, we did not have a "terminal" room in St. Louis – it was a full "computer room", with multiple PCs, workstations, X-windows terminals, laptop computer ethernet access, and printers.

Thanks to a cooperative arrangement with Southwest Bell, we had an innovative SMDS link at T1 speed. By the end of the week, there were some hard-working, but sleepy looking Wash University folks, who had spent some long hours providing these fine services. We owe them all a debt of thanks.

Of course, there is a downside that was pointed out by Vint Cerf. In the past, you could count on a low volume of email during IETF week, so that you wouldn't be too far behind after being out of the office for most of a week. At the last few IETF meetings, however, the connectivity has been so good that normal business seems to keep on going. For those of us who don't get to spend as much time in the "computer room" as we'd like, we are now assured of being buried in email by the end of the week!

This is a problem I'm pleased to have, so let me again thank Washington University, and our other recent hosts for presenting us with this unique "problem"!

Routing Protocols

There was a great deal of activity in the area of routing protocols at the March IETF meeting. The IESG routing Area Director, Bob Hinden (BBN), presented a set of criteria for advancing routing protocols through the Internet standardization process. Perhaps most far-reaching, was the notion that reports giving details of implementation experience and an analysis of the scaling behavior were requested. Asking about operational and implementation experience codifies the Internet philosophy that protocols should be tested before being standardized.

Two routing protocols were presented for advancment under these new criteria – OSPF and BGP. OSPF is an intra-AS routing protocol (i.e., IGP) which was proposed for advancement to Draft Standard. BGP is an inter-AS routing protocol, also proposed to advance to Draft Standard. Both protocols received a very thorough discussion, and the IESG expects to recommend both for advancement to Draft Standard.

There were some questions about the notion of a "common" IGP. By "common", we simply mean a protocol that is ubiquitously available from all router vendors (as in 'in common'). Users and network operators have expressed a strong need for routers from different vendors to have the capablity to interoperate within an AS through use of a common IGP.

In February 1990, the IESG recommended that the question of designating a "common" IGP be postponed until more information was available. More than a year has now passed since the IESG's recommendation, and many seemed to feel that it was now reasonable to re-open the consideration of designating a "common IGP". In particular, the Router Requirements document is now reaching closure, and many felt that a "common" IGP should designate in that document. The IESG agreed that a "common IGP" should be named in the Router Requirements document, and you can expect to see a recommendation from the IESG on this topic in time for the Atlanta IETF meeting.

More Kudos

Last, and certainly not least, I'd like to thank Megan Davies for coordinating the St. Louis Meeting and Candice Moshos for the outstanding job she did in assisting in that process. We are being spoiled with excellence at every turn!

Next IETF Meeting – Atlanta, July 29 - August 2

The next IETF meeting will be hosted by BellSouth and Caroline Cranfill in Atlanta (July 29 - August 2). Another good meeting is already taking shape. We hope to have joint sessions with the Trusted Systems Information Group (TSIG) to begin looking at a commercial IP security option and at ways to make NFS more secure, and, as mentioned above, we hope to open the issue of a "common IGP". I look forward to seeing you all there.

Final Agenda of the Twentieth IETF

(March 11-15, 1991)

MONDAY, March 11

8:00-9:00 am	IETF Registration and Continental Breakfast
9:00-9:30 am	Introductions and Local Arrangements
9:30-12:00 noon	Morning Sessions

- Interdomain Policy Routing WG (Martha Steenstrup/BBN)
- Internet Mail Extensions WG (Greg Vaudreuil/CNRI)
- Internet User Glossary WG (Karen Roubicek/BBN and Tracy LaQuey Parker/UTexas)
- Multi-Media Bridging WG (Jeffrey Fitzgerald/Fibercom)
- Security Area Advisory Group (observers welcome) (Stephen Crocker/TIS)
- Simple Network Management Protocol WG (Chuck Davin/MIT)
- Trouble Ticket Wishlist BOF (Dale Johnson/Merit)

10:45-11:00 am	Break (Coffee provided)
1:30-3:30 pm	Afternoon Sessions I

- Common Authentication Technology BOF (John Linn/DEC)
- FDDI MIB WG (Jeffrey Case/UTenn)
- Interdomain Policy Routing WG (Martha Steenstrup/BBN)
- Internet Mail Extensions WG (Greg Vaudreuil/CNRI)
- Introduction to Router Requirements (Philip Almquist/Consultant)
- IP over Appletalk WG (John Veizades/Apple)
- Multi-Media Bridging WG (Jeffrey Fitzgerald/Fibercom)
- Operational Statistics WG (Phill Gross/CNRI and Bernhard Stockman/NORDUnet)
- Remote LAN Monitoring WG (Mike Erlinger/Micro Technology)
- User Services WG* (Joyce Reynolds/ISI)

3:30-4:00 pm	Break (Refreshments provided)

4:00-6:00 pm — Afternoon Sessions II

- Internet Mail Extensions WG (Greg Vaudreuil/CNRI)
- IP over Appletalk WG (John Veizades/Apple)
- NOC-Tools WG (Robert Enger/Contel and Gary Malkin/FTP Software)
- Operational Statistics WG (Phill Gross/CNRI and Bernhard Stockman/NORDUnet)
- Password Management BOF (Jeffrey Schiller/MIT)
- Router Requirements WG (Philip Almquist/Consultant)

7:00-10:00 pm — Evening Sessions

- Conditioning of By-request Network Resources BOF (Andy Nicholson/CRAY)
- Management Services Interface WG (Oscar Newkerk/DEC)

TUESDAY, March 12, 1991

8:30-9:00 am	Continental Breakfast, No Morning Plenary
9:00-12:00 noon	Morning Sessions

- Connection IP WG (Claudio Topolcic/CNRI)
- Interdomain Policy Routing WG (Martha Steenstrup/BBN)
- Internet Accounting WG (Cyndi Mills/BBN)
- Internet Mail Extensions WG (Greg Vaudreuil/CNRI)
- Bridge MIB WG (Fred Baker/ACC)
- Router Requirements WG (Philip Almquist/Consultant)
- Network OSI Operational WG (Sue Hares/Merit)
- Security Policy WG (Rich Pethia/CERT)
- Telnet WG (Dave Borman/Cray Research)
- User Connectivity WG (Dan Long/BBN)
- X.400 Operations WG (Alf Hansen/UWisc)

10:30-10:45 am	Break (Coffee provided)
1:30-3:30 pm	Afternoon Sessions I

- Bridge MIB WG (Fred Baker/ACC)
- Connection IP WG (Claudio Topolcic/CNRI)
- Interdomain Policy Routing WG (Martha Steenstrup/BBN)
- Internet Accounting WG (Cyndi Mills/BBN)
- Internet Mail Extensions WG (Greg Vaudreuil/CNRI)
- Network Database WG (Daisy Shen/IBM)
- Operational Statistics WG (Phill Gross/CNRI and Bernhard Stockman/NORDUnet)
- Privacy Enhanced Mail I BOF (James Galvin/TIS)
- Service Location Protocol WG (John Veizades/Apple)
- Site Security Policy Handbook WG (Paul Holbrook/CICNet and Joyce Reynolds/ISI)
- X.400 Operations WG (Alf Hansen/UWisc)

3:30-4:00 pm	Break (Refreshments provided)
4:00-6:00 pm	Technical Presentations

- "Introduction to OSPF Simulation Demonstration" (Deepinder Sidhu/UMBC)
- "Architecture and Goals for the Interim Interagency National Research and Education Network (NREN)" (William Johnston/LBL and Peter Ford/LANL)

7:00-10:00 pm — Evening Sessions

- Conditioning of By-request Network Resources BOF (Andy Nicholson/CRAY)
- Multicast Extensions to OSPF WG (Steve Deering/Xerox PARC)
- Privacy Enhanced Mail II BOF (James Galvin/TIS)
- Remote LAN Monitoring WG (Mike Erlinger/Micro Technology)

WEDNESDAY, March 13, 1991

8:30-9:00 am	Continental Breakfast
9:00-9:30 am	Technical Presentations • "Routing Protocol Standardization Criteria" (Bob Hinden/BBN) • "DNS Issues" (Philip Almquist/Consultant)
9:30-12:00 noon	Morning Sessions • Connection IP WG (Claudio Topolcic/CNRI) • DECnet Phase IV MIB WG (Jonathan Saperia/DEC) • IP over Large Public Data Networks WG (George Clapp/Ameritech) • Interdomain Policy Routing WG (Martha Steenstrup/BBN) • Internet Accounting WG (Cyndi Mills/BBN) • IPSO/CIPSO BOF (Steve Crocker/TIS) • Network Printing Protocol WG (Glenn Trewitt/DEC) • Office Document Architecture WG (Peter Kirstein/UCL) • Router Requirements WG (Philip Almquist/Consultant) • Routing Simulator Demonstration (Deepinder Sidhu/UMBC) • Trouble Ticket Wishlist BOF (Dale Johnson/Merit) • X.400 Operations WG (Alf Hansen/UWisc)
10:30-10:45 am	Break (Coffee provided)
1:30-3:30 pm	Afternoon Sessions I • Border Gateway Protocol WG (Yakov Rekhter/IBM) • Character MIB WG (Bob Stewart/Xyplex) • Connection IP WG (Claudio Topolcic/CNRI) • Domain Name System WG (Michael Reilly/NSL) • IP over Large Public Data Networks WG (George Clapp/Ameritech) • Interdomain Policy Routing WG (Martha Steenstrup/BBN) • Internet Accounting WG (Cyndi Mills/BBN) • LISTSERV BOF (Phill Gross/CNRI) • Network Information Services Infrastructure WG (Dana Sitzler/Merit and Pat Smith/Merit) • Network Joint Management WG (Gene Hastings/PSU) • Telnet WG (Dave Borman/Cray Research) • X.400 Operations WG (Alf Hansen/UWisc) • Routing Simulator Demonstration (Deepinder Sidhu/UMBC)
3:30-4:00 pm	Break (Refreshments provided)

4:00-6:00 pm	Technical Presentions	

- "NREN Legislative Update" (Bill Bostwick/Federal Networking Council)
- "Advanced Network Research at Washington University" (Jonathan Turner/WashU)
- "Experiments with DEC-bit Congestion Avoidance" (Rick Wilder/Mitre)

7:00-10:00pm Evening Session

- Border Gateway Protocol WG (Yakov Rekhter/IBM)
- SNMP Authentication WG (James Galvin/TIS)

THURSDAY, March 14, 1991

8:30-9:00 am	Continental Breakfast, No Morning Plenary
9:00-12:00 noon	Morning Sessions

- Border Gateway Protocol WG* (Yakov Rekhter/IBM)
- Connection IP WG (Claudio Topolcic/CNRI)
- Dynamic Host Configuration WG (Ralph Droms/Bucknell)
- Internet Accounting WG (Cyndi Mills/BBN)
- IP over Large Public Data Networks WG* (George Clapp/Ameritech)
- Network Fax WG (Mark Needleman/UC)
- Open Shortest Path First IGP WG (John Moy/Proteon)
- Router Requirements WG (Philip Almquist/Consultant)
- Security Area Advisory Group (observers welcome) (Stephen Crocker/TIS)
- Network Status Reports (Phill Gross/CNRI)

10:30-10:45 am	Break (Coffee provided)
1:30-4:00 pm	Technical Presentations

- "OSPF Status Report" (John Moy/Proteon, Jeff Burgan/NASA and Rob Coltun/Wellfleet)
- "BGP Status Report" (Yakov Rekhter/IBM Dennis Ferguson/CA*net and Sue Hares/Merit)
- "Alert Management Status Report " (Louis Steinberg/IBM)
- "IETF Internet Access Via SMDS" (Kirk Williams/Southwestern Bell Technology Resources)

4:00-4:30 pm	Break (Refreshments provided)
4:30-6:00 pm	Open Plenary and IESG
7:00pm-10:00	• A Tutorial on Office Document Architecture (Peter Kirstein/UCL)

* BGP and IPLPDN will meet jointly from 9:00-10:30. After the "Break" they will meet separately.

FRIDAY, March 15, 1991

8:30-9:00 am	Continental Breakfast
9:00-11:30 am	Working Group Area and Selected Working Group Presentations

- "Introduction to Frame Relay" (Andrew Malis/BBN)
- User Services Area (Joyce K. Reynolds/ISI)
- Applications Area (Russ Hobby/UC Davis)
- Internet Services Area (Noel Chiappa/Consultant)
- Routing Area (Bob Hinden/BBN)
- Security Area (Steve Crocker/TIS)
- OSI Integration Area (Ross Callon/DEC and Rob Hagens/UWisc)
- Operations Area (Interim - Phill Gross/CNRI)
- Network Management Area (Chuck Davin/MIT)

10:15-10:30 am	Break (Coffee provided)
11:30-12:00 noon	Concluding Remarks (Phill Gross/CNRI)
12:00 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 the 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. Together with the Chair of the IETF, these eight technical Directors compose the Internet Engineering Steering Group (IESG).

The current Areas and Directors, which compose the IESG, are:

IETF and IESG Chair:	Phill Gross/CNRI
Applications:	Russ Hobby/UC-Davis
Internet and	Noel Chiappa/Consultant
Transport Services	Dave Borman/Cray
Network Management:	James Davin/ MIT
OSI Integration:	Rob Hagens/UWisc and Ross Callon/DEC
Operational Requirements:	Phill Gross/CNRI (interim)
Routing:	Robert Hinden/BBN
Security:	Steve Crocker/TIS
User Services	Joyce Reynolds/ISI
Standards Management:	Dave Crocker/DEC
IESG Secretary:	Greg Vaudreuil/CNRI

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 plenary sessions three times a year composed of Working Group Sessions, Technical Presentations and Network Status Briefings. The meetings are currently four and one half days long and include 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.

Mailing Lists

Much of the daily work of the IETF is conducted on electronic mailing lists. There are mailing lists for each of the Working Groups, as well as a general IETF list. Mail on the Working Group mailing lists is expected to be technically relevant to the Working Groups supported by that list.

To join a mailing list, send a request to the associated request list. All internet mailing lists have a companion "-request" list. Send requests to join a list to <listname>-request@<listhost>.

Information and logistics about upcoming meetings of the IETF are distributed on the general IETF mailing list. For general inquiries about the IETF, send a request to `ietf-request@isi.edu`. An archive of mail sent to the IETF list is available for anonymous ftp from the directory `~ftp/irg/ietf` on `venera.isi.edu`

1.1 Future IETF Meeting Sites

Summer 1991

BellSouth Services
Host: Caroline Cranfill
July 29th - August 2nd

Fall/ Winter 1991

Los Alamos National Laboratory
Host: John Morrison and Dale Land
November 18th - 22nd

Spring 1992

San Diego Supercomputer Center
Host: E. Paul Love, Jr. and Hans-Werner Braun
March 92 (tentative)

1.2 On Line IETF Information

The Internet Engineering Task Force maintains up-to-date, on-line information on all its activities. This information is available via FTP through the NSFnet Service Center (NNSC) and through several "shadow" machines. These "shadow" machines may in fact be more convenient than the NNSC. Procedures for retrieving the information are listed below.

Directory Locations

Information pertaining to the IETF, its Working Groups and Internet Drafts can be found in either the "IETF" Directory or the "Internet-Drafts" Directory. (For a more detailed description of these Directories, please see Section 1.2.1 and 1.2.2). To retrieve this information via FTP, establish a connection, then Login with username ANONYMOUS and password GUEST. When logged in, change to the directory of your choice with the following commands:

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

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 (192.31.103.6) **The Defense Data Network NIC** Address: nic.ddn.mil (192.67.67.20)

> 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 6:00am and 5:00pm Pacific time.

Pacific Rim Address: munnari.oz.au (128.250.1.21)

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

Europe Address: nic.nordu.net (192.36.148.17)

- This machine will accept only an email address as the password.

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 IETF plenary, containing scheduled Working Groups meetings, Technical Presentations and Network Status Reports.
0mtg-at-a-glance	The announcement for the upcoming IETF plenary, containing specific information on the date/location of the meeting, hotel/airline arrangements, meeting site accommodations and meeting costs.
0mtg-rsvp	A standardized RSVP form to notify the staff of your plans to attend the upcoming IETF meeting.
0mtg-sites	Current and future meeting dates and sites for IETF plenaries.
1id-abstracts	The Internet Drafts currently on-line in the Internet-Drafts directory.
1id-guidelines	Instructions for authors of Internet Drafts.
1ietf-description	A short description of the IETF, the IESG and how to participate.
1wg-summary	A listing of all current Working Groups, the Working Group Chairs 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 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 RFC's. These documents are indexed in the file 1id-abstracts.txt in the Internet-Drafts directory. 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.

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".

1.3 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 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. We do however, strongly encourage 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 Internet Draft 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 an RFC or a Draft RFC.

The Internet Draft should neither 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

Distributed Scheduling Protocol (chronos)
Chair(s): Paul Linder lindner@boombox.micro.umn.edu
WG mail: chronos@boombox.micro.umn.edu
To Join: chronos-request@boombox.micro.umn.edu

Internet Mail Extensions (smtpext)
Chair(s): Gregory Vaudreuil gvaudre@nri.reston.va.us
WG mail: ietf-smtp@dimacs.rutgers.edu
To Join: ietf-smtp-request@dimacs.rutgers.edu

Internet Message Extentions (822ext)
Chair(s): Gregory Vaudreuil gvaudre@nri.reston.va.us
WG mail: ietf-822@dimacs.rutgers.edu
To Join: ietf-822-request@dimacs.rutgers.edu

Network Database (netdata)
Chair(s): Daisy Shen daisy@watson.ibm.com
WG mail: ietf-ndb@ucdavis.edu
To Join: ietf-ndb-request@ucdavis.edu

Network News Transport Protocol (nntp)
Chair(s): Eliot Lear
WG mail: ietf-nntp@turbo.bio.net
To Join: ietf-nntp-request@turbo.bio.net

Distributed File Systems (dfs)
Chair(s): Peter Honeyman honey@citi.umich.edu
WG mail: dfs-wg@citi.umich.edu
To Join: dfs-wg-request@citi.umich.edu

Domain Name System (dns)
Chair(s): Michael Reilly reilly@pa.dec.com
WG mail: dns-wg@nsl.dec.com
To Join: dns-wg-request@nsl.dec.com

Network Fax (netfax)
Chair(s): Mark Needleman mhn@stubbs.ucop.edu
WG mail: netfax@stubbs.ucop.edu
To Join: netfax-request@stubbs.ucop.edu

Network Printing Protocol (npp)
Chair(s): Glenn Trewitt trewitt@pa.dec.com
WG mail: print-wg@pluto.dss.com
To Join: print-wg-request@pluto.dss.com

TELNET (telnet)
Chair(s): Dave Borman dab@cray.com
WG mail: telnet-ietf@cray.com
To Join: telnet-ietf-request@cray.com

Network Management

J.R. Davin
jrd@ptt.lcs.mit.edu

Remote LAN Monitoring (rlanmib)
Chair(s): Mike Erlinger mike@mti.com
WG mail: rlanmib@mti.com
To Join: rlanmib-request@mti.com

X.25 Management Information Base (x25mib)
Chair(s): Dean Throop throop@dg-rtp.dg.com
WG mail: x25mib@dg-rtp.dg.com
To Join: x25mib-request@dg-rtp.dg.com

Bridge MIB (bridge)
Chair(s): Fred Baker fbaker@emerald.acc.com
WG mail: bridge-mib@nsl.dec.com
To Join: bridge-mib-request@nsl.dec.com

Character MIB (charmib)
Chair(s): Bob Stewart rlstewart@eng.xyplex.com
WG mail: char-mib@decwrl.dec.com
To Join: char-mib-request@decwrl.dec.com

DECnet Phase IV MIB (decnetiv)
Chair(s): Jonathan Saperia saperia@decwrl.enet.dec.com
WG mail: phiv-mib@jove.pa.dec.com
To Join: phiv-mib-request@jove.pa.dec.com

FDDI MIB (fddimib)
Chair(s): Jeffrey Case case@cs.utk.edu
WG mail: fddi-mib@CS.UTK.EDU
To Join: fddi-mib-request@CS.UTK.EDU

Internet Accounting (acct)
Chair(s): Cyndi Mills cmills@bbn.com
WG mail: accounting-wg@bbn.com
To Join: accounting-wg-request@bbn.com

Management Services Interface (msi)

Chair(s):	Oscar Newkerk	newkerk@decwet.enet.dec.com
	Sudhanshu Verma	verma@hpindbu.cup.hp.com
WG mail:	msiwg@decwrl.dec.com	
To Join:	msiwg-request@decwrl.dec.com	

OSI Internet Management (oim)

Chair(s):	Lee LaBarre	cel@mbunix.mitre.org
	Brian Handspicker	bd@vines.enet.dec.com
WG mail:	oim@mbunix.mitre.org	
To Join:	oim-request@mbunix.mitre.org	

SNMP Security (snmpsec)

Chair(s):	James Galvin	galvin@tis.com
	Keith McCloghrie	kzm@hls.com
WG mail:	snmp-sec-dev@tis.com	
To Join:	snmp-sec-dev-request@tis.com	

Simple Network Management Protocol (snmp)

Chair(s):	Marshall Rose	mrose@psi.com
WG mail:	snmp-wg@nisc.nyser.net	
To Join:	snmp-wg-request@nisc.nyser.net	

OSI Integration

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

Network OSI Operations (noop)
Chair(s): Susan Hares skh@merit.edu
WG mail: noop@merit.edu
To Join: noop-request@merit.edu

Office Document Architecture (oda)
Chair(s): Peter Kirstein kirstein@cs.ucl.ac.uk
WG mail: ietf-osi-oda@cs.ucl.ac.uk
To Join: ietf-osi-oda-request@cs.ucl.ac.uk

X.400 Operations (x400ops)
Chair(s): Alf Hansen Alf.Hansen@pilot.cs.wisc.edu
WG mail: ietf-osi-x400ops@pilot.cs.wisc.edu
To Join: ietf-osi-x400ops-request@pilot.cs.wisc.edu

Assignment of OSI NSAP Addresses (osinsap)
Chair(s): Richard Colella colella@osi3.ncsl.nist.gov
WG mail: ietf-osi-nsap@osi3.ncsl.nist.gov
To Join: ietf-osi-nsap-request@osi3.ncsl.nist.gov

OSI Directory Services (osids)
Chair(s): Steve Kille S.Kille@cs.ucl.ac.uk
WG mail: ietf-osi-ds@cs.ucl.ac.uk
To Join: ietf-osi-ds-request@cs.ucl.ac.uk

OSI General (osigen)
Chair(s): Robert Hagens hagens@cs.wisc.edu
Ross Callon callon@bigfut.enet.dec.com
WG mail: ietf-osi@cs.wisc.edu
To Join: ietf-osi-request@cs.wisc.edu

OSI X.400 (osix400)
Chair(s): Rob Hagens hagens@cs.wisc.edu
WG mail: ietf-osi-x400@cs.wisc.edu
To Join: ietf-osi-x400-request@cs.wisc.edu

Operational Requirements

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

DDN Interconnectivity (ddniwg)
Chair(s): Kathleen Huber khuber@bbn.com
WG mail:
To Join:

Operational Statistics (opstat)
Chair(s): Bernhard Stockman bygg@sunet.se
Phillip Gross pgross@nri.reston.va.us
WG mail:
To Join:

Benchmarking Methodology (bmwg)
Chair(s): Scott Bradner sob@harvard.edu
WG mail: bmwg@harvisr.harvard.edu
To Join: bmwg-request@harvisr.harvard.edu

Network Joint Management (njm)
Chair(s): Gene Hastings hastings@psc.edu
WG mail: njm@merit.edu
To Join: njm-request@merit.edu

Topology Engineering (tewg)
Chair(s): Not Yet Filled
WG mail: tewg@devvax.tn.cornell.edu
To Join: tewg-request@devvax.tn.cornell.edu

User Connectivity (ucp)
Chair(s): Dan Long long@nic.near.net
WG mail: ucp@nic.near.net
To Join: ucp-request@nic.near.net

Routing

Bob Hinden
hinden@bbn.com

Border Gateway Protocol (bgp)
Chair(s): Yakov Rekhter yakov@ibm.com
WG mail: iwg@rice.edu
To Join: iwg-request@rice.edu

IP over Large Public Data Networks (iplpdn)
Chair(s): George Clapp meritec!clapp@uunet.uu.net
WG mail: iplpdn@nri.reston.va.us
To Join: iplpdn-request@nri.reston.va.us

ISIS for IP Internets (isis)
Chair(s): Ross Callon callon@bigfut.enet.dec.com
WG mail: isis@merit.edu
To Join: isis-request@merit.edu

Inter-Domain Policy Routing (idpr)
Chair(s): Martha Steenstrup msteenst@bbn.com
WG mail: idpr-wg@bbn.com
To Join: idpr-wg-request@bbn.com

Multicast Extentions to OSPF (mospf)
Chair(s): Steve Deering deering@xerox.com
WG mail: mospf@devvax.tn.cornell.edu
To Join: mospf-request@devvax.tn.cornell.edu

Open Shortest Path First IGP (ospf)
Chair(s): Mike Petry petry@trantor.umd.edu
John Moy jmoy@proteon.com
WG mail: ospfigp@trantor.umd.edu
To Join: ospfigp-request@trantor.umd.edu

Security

Steve Crocker
crocker@tis.com

Commercial Internet Protocol Security Option (cipso)

Chair(s):	Ron Sharp	rls@neptune.att.com
WG mail:	cipso@wdl1.wdl.loral.com	
To Join:	cipso-request@wdl1.wdl.loral.com	

Common Authentication Technology (cat)

Chair(s):	John Linn	zendia.enet.dec.com
WG mail:	TBD	
To Join:		

IP Authentication (ipauth)

Chair(s):	Jeffrey Schiller	jis@mit.edu
WG mail:	awg@bitsy.mit.edu	
To Join:	awg-request@bitsy.mit.edu	

Internet Security Policy (spwg)

Chair(s):	Richard Pethia	rdp@cert.sei.cmu.edu
WG mail:	spwg@nri.reston.va.us	
To Join:	spwg-request@nri.reston.va.us	

Site Security Policy Handbook (ssphwg)

Chair(s):	J. Paul Holbrook	holbrook@cic.net
	Joyce K. Reynolds	jkrey@isi.edu
WG mail:	ssphwg@cert.sei.cmu.edu	
To Join:	ssphwg-request@cert.sei.cmu.edu	

User Services

Joyce Reynolds
jkrey@isi.edu

Directory Information Services Infrastructure (disi)

Chair(s):	Chris Weider	clw@merit.edu
WG mail:	disi@merit.edu	
To Join:	disi-request@merit.edu	

Internet User Glossary (userglos)

Chair(s):	Karen Roubicek	roubicek@bbn.com
	Tracy Parker	tracy@utexas.edu
WG mail:	usergloss@ftp.com	
To Join:	usergloss-request@ftp.com	

NOC-Tool Catalogue Revisions (noctool2)

Chair(s):	Robert Enger	enger@seka.scc.com
	Gary Malkin	gmalkin@ftp.com
WG mail:	noctools@merit.edu	
To Join:	noctools-request@merit.edu	

Network Information Services Infrastructure (nisi)

Chair(s):	Dana Sitzler	dds@merit.edu
	Pat Smith	psmith@merit.edu
WG mail:	nisi@merit.edu	
To Join:	nisi-request@merit.edu	

User Services (uswg)

Chair(s):	Joyce K. Reynolds	jkrey@isi.edu
WG mail:	us-wg@nnsc.nsf.net	
To Join:	us-wg-request@nnsc.nsf.net	

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. These drafts are listed alphabetically by Working Group acronym and initial post date. Drafts not originating in a Working Group are listed first.

"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.

"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.

"Asynchronous Discovery of an Effective Maximum Transmission Unit for IP Datagram Delivery [MTU Discovery]", James Sawyer, 08/17/1990 <draft-csc-sawyer-mtudisc-00.txt>

> A case against IP layer fragmentation has been made, and various methods for avoiding it proposed. This memo revisits the effect of fragmentation on network performance, and recounts the present methods of avoidance. A protocol is presented which adapts to the varying circumstances encountered, sending large datagrams whenever possible, and reducing fragmentation when necessary to avoid retransmission problems. A hybrid approach to MTU discovery, it utilizes one new IP header option and four new ICMP messages. It is a simple mechanism that discovers path MTUs without wasting resources and that works well before all hosts and routers are modified.

"FTP-FTAM Gateway Specification", J.L. Mindel, R.L. Slaski, 11/19/1990 <draft-slaski-ftpftam-00.txt>

> This memo describes a dual protocol stack application layer gateway that performs protocol translation, in an interactive environment, between the FTP and FTAM file transfer protocols. Only through additional implementations and fieldings will the FTP-FTAM gateway reach its optimal capacity as a resource during the anticipated long coexistence of the TCP/IP and OSI protocol suites.

"Tunneling IPX Traffic through IP Networks", Don Provan, 04/15/1991 <draft-provan-ipxtunneling-01.txt>

Internet Packet eXchange protocol (IPX) is the internetwork protocol used by Novell's NetWare protocol suite. For the purposes of this paper, IPX is functionally equivalent to the Internet Datagram Protocol (IDP) from the Xerox Network Systems (XNS) protocol suite. This draft specifies a method of encapsulating IPX datagrams within UDP packets so that IPX traffic can travel across an IP internet. This draft allows an IPX implementation to view an IP internet as a single IPX network. An implementation of this draft will encapsulate IPX datagrams in UDP packets in the same way any hardware implementation might encapsulate IPX datagrams in that hardware's frames. IPX networks can be connected thusly across internets that carry only IP traffic.

"U.S. Department of Defense Security Options for the Internet Protocol", Stephen Kent, 03/05/1991 <draft-ietf-ahwgipso-ipso-00.txt>

This I-D specifies the U.S. Department of Defense Basic Security Option and the top-level description of the Extended Security Option for use with the Internet Protocol. This proposal will replace RFC 1038 Revised "IP Security Option", dated January 1988, if it is issued as an RFC.

"The IP Addressing Issue", Noel Chiappa, 03/27/1991 <draft-chiappa-ipaddressing-00.txt>

The packet layer of the IP architecture is about to enter a period of stress caused by deficiencies in the IP address. This stress is caused by a number of inter-related problems. This note describes these problems, lists some suggested solutions, and discusses pros and cons of each of those solutions.

"Internet Accounting: Background", C. Mills, D. Hirsh, G. Ruth, , 05/13/1991 <draft-ietf-acct-background-00.txt>

This document provides background information for the "Internet Accounting Architecture" and is the first of a three document set: Internet Accounting Background & Status (this document), Internet Accounting Architecture (under construction), Internet Accounting Meter Service (under construction). The focus at this time is on defining METER SERVICES and USAGE REPORTING which provide basic semantics for measuring network utilization, a syntax, and a data reporting protocol. The intent is to produce a set of standards that is of practical use for early experimentation with usage reporting as an internet accounting mechanism. This document provides background and tutorial information on issues surrounding the architecture, or in a sense, an explanation of choices made in the Internet Accounting Architecture.

"AppleTalk MIB", Steven Waldbusser, 02/11/1991 <draft-ietf-appleip-applemib-00.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 defines objects for managing AppleTalk networks.

"The Transmission of IP Datagrams Over AppleTalk Networks", John Veizades, 03/08/1991 <draft-ietf-appleip-ipoverappletalk-00.txt>

This document describes a protocol, called MacIP, that is used to transport IP datagrams on AppleTalk networks. This protocol was developed in order to connect Macintosh computers on AppleTalk networks to hosts on TCP/IP networks. Using the AppleTalk network layer protocol, IP datagrams can be transmitted through AppleTalk networks to gateways that decapsulate the IP datagrams and act as front-end protocol processors Macintosh hosts on AppleTalk internets.

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

This draft 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.

"Definitions of Managed Objects for the Border Gateway Protocol (Version 3)", Steven Willis, John Burruss, 03/25/1991 <draft-ietf-iwg-bgp-mib-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 defines objects for managing the Border Gateway Protocol.

"A Border Gateway Protocol 3 (BGP-3)", Yakov Rekhter, Kirk Lougheed, 01/25/1991 <draft-ietf-bgp-bgp3-00.txt>

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 and EGP usage in the NSFNET Backbone as described in RFC 1092 and RFC 1093. The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems. This networkreachability information includes information on the full path of 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 some policy decisions at the AS level may be enforced.

"Border Gateway Protocol NEXT-HOP-SNPA Attribute", Paul Tsuchiya, 04/15/1991 <draft-ietf-bgp-nexthop-00.txt>

The purpose of the NEXT-HOP attribute is for one border gateway A to tell another border gateway B which border gateway C (where C might equal A) should be used as the next hop on the path to the destinations advertised in the UPDATE containing the NEXT-HOP attribute.

"Experience with the BGP Protocol", Yakov Rekhter, 05/08/1991 <draft-ietf-bgp-experience-00.txt>

The purpose of this memo is to document how the requirements for advancing a routing protocol to Draft Standard have been satisfied by BGP. This report documents experience with BGP.

"BGP Protocol Analysis", Yakov Rekhter, 05/08/1991 <draft-ietf-bgp-analysis-00.txt>

The purpose of this report is to document how the requirements for advancing a routing protocol to Draft Standard have been satisfied by BGP. This report summarizes the key feature of BGP, and analyzes the protocol with respect to scaling and performance.

"Benchmarking Terminology for Network Interconnection Devices", Scott Bradner, 11/26/1990 <draft-ietf-bmwg-terms-01.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.

"Definitions of Managed Objects for Bridges", E. Decker, P. Langille,, A. Rijsinghani, K. McCloghrie, 05/24/1991 <draft-ietf-bridge-definitions-01.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 defines objects for managing bridges based on the IEEE 802.1d draft standard between Local Area Network (LAN) segments. Provisions are made for support of transparent and source route bridging. Provisions are also made so that these objects apply to bridges connected by subnetworks other than LAN segments.

"Definitions of Managed Objects for Parallel-printer-like Hardware Devices", Bob Stewart, 01/02/1991 <draft-ietf-charmib-parallelprinter-01.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 defines objects for managing parallel-printer-like hardware devices.

"Definitions of Managed Objects for RS-232-like Hardware Devices", Bob Stewart, 01/02/1991 <draft-ietf-charmib-rs232like-01.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 defines objects for managing RS-232-like hardware devices.

"Definitions of Managed Objects for Character Stream Devices", Bob Stewart, 01/02/1991 <draft-ietf-charmib-charmib-01.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 defines objects for managing character stream devices.

"DECnet Phase IV MIB Extensions", Jon Saperia, 06/06/1991 <draft-ietf-decnetiv-mibextensions-00.txt, .ps>

This memo defines a set of DECnet Phase IV extensions that have been created for the Internet MIB. When used in conjunction with the structure of management information (RFC 1155), the management information base for network management of TCP/IP-based internets (RFC 1156) and the Simple Network Management Protocol (RFC 1157), it will be possible to provide integrated network management of combined TCP/IP and DECnet Phase IV based internets. This document was produced by the DECnet Phase IV MIB working group.

"Clarifications and Extensions for the Bootstrap Protocol", Walt Wimer, 05/03/1991 <draft-ietf-dhc-bootp-00.txt>

Some aspects of the BOOTP protocol were rather loosely-defined in its original specification. In particular, only a general description was provided for the behavior of "BOOTP relay agents" (originally called "BOOTP forwarding agents"). The client behavior description also suffered in certain ways. This memo attempts to clarify and strengthenthe specification in these areas.

In addition, new issues have arisen since the original specification was written. This memo also attempts to address some of these.

"INTERNET OSI INTEGRATION, COEXISTENCE AND INTEROPERABILITY ISSUES", Robert Hagens, Rebecca Nitzan, 07/24/1990 <draft-fopg-ositransition-00.txt>

The intent of this document is to provide technical descriptions of the issues involved in the integration of the Open Systems Interconnect (OSI) protocols into the operational networks which interconnect and comprise the "Internet". The issues raised and solutions discussed are a result of the Federal Networking Council (FNC) OSI Planning Group (FOPG). The members of the FOPG represent several Federal Government agencies such as the Department of Energy (DOE), the National Science Foundation (NSF) the National Aeronautics and Space Administration (NASA), the National Institute of Standards and Technology (NIST) under the Department of Commerce, as well as University experts.

"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 transitted. The architecture is designed to accommodate an Internet with tens of thousands of administrative domains.

"Inter-Domain Policy Routing Protocol Specification and Usage: Version 1", M. Steenstrup, 03/05/1991 <draft-ietf-idpr-specv1-00.txt, or .ps>

We present the version 1 protocols for inter-domain policy routing (IDPR), which include the virtual gateway protocol, the domain status distribution protocol, the route synthesis procedure, the path setup protocol, and the message forwarding procedure. We also supply protocol usage recommendations to simplify implementation of IDPR.

"Internet Routing Protocol Standardization Criteria", Bob Hinden, 04/01/1991 <draft-ietf-iesg-routing-00.txt>

The purpose of this document is to provide more specific guidance for the advancement of routing protocols. All levels of the standardization process are covered.

"Management Services Application Programming Interface", Oscar Newkerk, 12/12/1990 <draft-ietf-msi-api-03.txt, or .ps>

A case against IP layer fragmentation has been made, and various methods for avoiding it proposed. This memo revisits the effect of fragmentation on network performance, and recounts the present methods of avoidance. A protocol is presented which adapts to the varying circumstances encountered, sending large datagrams whenever possible, and reducing fragmentation when necessary to avoid retransmission problems. A hybrid approach to MTU discovery, it utilizes one new IP header option and four new ICMP messages. It is a simple mechanism that discovers path MTUs without wasting resources and that works well before all hosts and routers are modified.

"Network Time Protocol: Version 3", Dave Mills, 11/28/1990 <draft-mills-ntpv3-00.txt, or .ps>

This document describes the Network Time Protocol (NTP), specifies its formal structure and summarizes information useful for its implementation. NTP provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse internet operating at rates from mundane to lightwave. It uses a returnable-time design in which a distributed subnet of time servers operating in a self- organizing, hierarchical-master-slave configuration synchronizes local clocks within the subnet and to national time standards via wire or radio. The servers can also redistribute reference time via local routing algorithms and time daemons.

"The IP Network Address Translator (Nat): Preliminary Design", Paul Tsuchiya, 04/15/1991 <draft-tsuchiya-addrtrans-00.txt, .ps>

The two most compelling problems facing the IP Internet are IP address depletion and scaling in routing. This paper discusses the characteristics of one of the proposed solutions-address reuse. The solution is to place Network Address Translators (Nat) at the borders of stub domains. Each Nat box has a small pool of globally unique IP addresses that are dynamically assigned to IP flows going through Nat. The dynamic assignment is coordinated with the Domain Name Servers. The IP addresses inside the stub domain are not globally unique-they are reused in other domains, thus solving the address depletion problem. The pool of IP addresses in Nat is from a subnet administered by the regional backbone, thus solving the scaling problem. The main advantage of Nat is that it can be installed without changes to routers or hosts. This paper presents a preliminary design for Nat, and discusses its pros and cons.

"WORKSHOP ON CO/CL INTERWORKING", Phill Gross, Les Clyne, COCL Workshop, , 12/12/1990 <draft-ccirn-cocl-report-00.txt>

On July 24-26, 1990, an invited panel met at the Corporation for National Research Initiatives in Reston Virginia to consider the issues involved with interworking between protocol stacks based on Connection-mode Network Service (CONS, or CO) and Connectionless-mode Network Service (CLNS, or CL). The main example of a CO stack is OSI TP0 over X.25. Examples of CL protocolstacks include OSI TP4 over CLNP and TCP over IP. The workshop was convened at the direction of RARE and the U.S. Federal Networking Council (FNC). The meeting was organized and co-chaired by Les Clyne (UK Joint Network Team) and Phillip Gross(Corporation for National Research Initiatives). An electronic mailing list was established for use by both attendees and a wider audience of experts. This report gives an overview and synopsis of the deliberations at the meeting, and it describes the outcome.

"An Approach to CO/CL Interworking – Part I: Introduction", COCL Workshop, M. Rose, 05/06/1991 <draft-ccirn-cocl-doc1-02.txt>

The OSI transport service may be realized through a variety of transport/network protocol combinations. Regrettably, few of the combinations actually interoperate with each other. As such, even if all OSI-capable end-systems enjoyed full-connectivity, they would not be able to uniformly interoperate. This memo examines the problem and proposes an approach in order to develop solutions to this problem.

"An Approach to CO/CL Interworking– Part II: The Short-Term – Conventions for Transport-Service Bridges in the Absence of Internetworking", COCL Workshop, M Rose, 04/23/1991 <draft-ccirn-cocl-doc2-01.txt>

The Short-term approach outlined in "An Approach to CO/CL Interworking: Part I: Introduction" is based on the use of transport-layer relays known as transport service bridges, or TS-bridges. Further, the short-term approach also assumes that knowledge of the TS-bridges is present in the end-systems. The companion memo "An Approach to CO/CL Interworking–Part III: The Intermediate-Term–Provision of the CONS over TCP and X.25 Subnetworks" identifies solutions in which end-system knowledge of transport-layer relays is avoided. The purpose of this memo is two-fold: first, modifications to the operational characteristics of end-systems are described; and, second, the operational characteristics of TS-bridges are described.

"An Approach to CO/CL Interworking – Part III: The Long-Term – Conventions for Network-Layer Relays and Transport-Service Bridges in the presence of Internetworking", CO/CL Workshop, C. Huitema, 04/25/1991 <draft-ccirn-cocl-doc4-01.txt>

The long-term approach is based on the use of transport-layer relays known as transport service bridges, or TS-bridges. Further, the long-term approach also assumes that knowledge of the TS-bridges is hidden from the end-systems. The companion memo identifies the short-term approach towards TS-bridges. The purpose of this memo is three-fold: first, to identify the infrastructure which is expected to exist in the long-term; second, to describe the use of NL-relays in such an environment. And, third, to describe the use of TS-bridges in such an environment.

"Mapping between X.400(1988) / ISO 10021 and RFC 822", S.E. Hardcastle-Kille, 05/31/1991 <draft-ietf-kille-x.400mapping-00.txt>

This document describes a set of mappings which will enable interworking between systems operating the CCITT X.400 (1988) Recommendations on Message Handling Systems/ISO IEC 10021 Message Oriented Text Interchange Systems (MOTIS), and systems using the RFC 822 mail protocol or protocols derived from RFC 822. The approach aims to maximise the services offered across the boundary, whilst not requiring unduly complex mappings. The mappings should not require any changes to end systems. This document is a revision based on RFCs 987, 1026, 1138, and 1148 which it obsoletes.

"A String Encoding of Presentation Address", S.E. Kille, 01/16/1991 <draft-ucl-kille-presentationaddress-02.txt, or .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, 03/21/1991 <draft-ucl-kille-x500domains-03.txt, or .ps>

This draft 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 emphasized. In addition, a mapping of domains onto X.500 is proposed, which gives a range of new management and user facilities over and above those currently available. This specification proposes an experimental new mechanism to access and manage domain information on the Internet and in the UK Academic Community. There is no current intention to provide an operational replacement for DNS. Please send comments to the author or to the discussion group osi-ds@CS.UCL.AC.UK.

"An Interim Approach to use of Network Addresses", S. Kille, 01/14/1991 <draft-ucl-kille-networkaddresses-02.txt, or .ps>

This note is a proposal for mechanisms to utilize Network Addresses. 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.

"Replication Requirement to Provide an Internet Directory Using X.500", S. Kille, 03/21/1991 <draft-ietf-osids-replication-02.txt, or .ps>

A companion document discussed an overall framework for deploying X.500 on the Internet "Building and internet directory using X.500" . This document considers certain deficiencies of the 1988 standard, which need to be addressed before an effective open Internet Directory can be established. The only areas considered are primary problems, to which solutions must be found before a pilot can be deployed. This INTERNET-DRAFT concerns itself with deficiencies which can only be addressed by use of additional protocol or procedures for distributed operation.

"Using the OSI Directory to Achieve User Friendly Naming", S. Kille, 03/21/1991 <draft-ietf-osids-friendlynaming-02.txt, or .ps>

This proposal sets out some conventions for representing names in a friendly manner, and shows how this can be used to achieve really friendly naming. This then leads to a specification of a standard format for representing names, and to procedures to resolve them. Please send comments to the author or to the discussion group osi-ds@CS.UCL.AC.UK.

"Replication and Distributed Operations Extensions to Provide an Internet Directory using X.500", S. Kille, 03/21/1991 <draft-ietf-osids-replsoln-02.txt, or .ps>

Some requirements on extensions to X.500 are described in the INTERNET DRAFT "Replication requirement to provide an internet, in order to build an Internet Directory Using X.500", as described in the INTERNET DRAFT "Building and internet directory using X.500". This document specifies a set of solutions to the problems raised. These solutions are based on some work done for the QUIPU implementation, and demonstrated to be effective in a number of directory pilots. By documenting a de facto standard, rapid progress can be made towards a full-scale pilot. These procedures are an INTERIM approach. There are known deficiencies, both in terms of manageability and scalability. Transition to standard approaches are planned when appropriate standards are available. This INTERNET DRAFT will be obsoleted at this point.

"The COSINE and Internet X.500 Schema", P. Barker, S. Kille, 03/21/1991 <draft-ietf-osids-cosinex500-03.txt>

This document suggests an X.500 Directory Schema, or Naming Architecture, for use in the COSINE and Internet X.500 pilots. The schema is independent of any specific implementation. As well as indicating support for the standard object classes and attributes, a large number of generally useful object classes and attributes are also defined. An appendix to this document includes a machine processable version of the schema. This document also proposes a mechanism for allowing the schema to evolve in line with commonly held requirements. Proform as to support this process are included. Please send comments to the authors or to the discussion group osi-ds@cs.ucl.ac.uk.

"Handling QOS (Quality of service) in the Directory", S.E. Kille, 03/20/1991 <draft-ietf-osids-qos-00.txt, or .ps>

This document describes a mechanism for specifying the Quality of Service for DSA Operations and Data in the Internet Pilot Directory Service "Building and internet directory using X.500". Please send comments to the author or to the discussion group osi-ds@CS.UCL.AC.UK.

"DSA Naming", S.E. Kille, 03/21/1991 <draft-ietf-osids-dsanaming-00.txt, or .ps>

This INTERNET-DRAFT describes a few problems with DSA Naming as currently deployed in pilot exercises, and suggests a new approach. This approach is suggested for use in the Internet Directory Pilot. Please send comments to the author or to the discussion group osi-ds@CS.UCL.AC.UK.

"Naming Guidelines for Directory Pilots", P. Barker, S.E. Kille, 03/21/1991 <draft-ietf-osids-dirpilots-00.txt, or .ps>

Deployment of a Directory will benefit from following certain guidelines. This document defines a number of guidelines which are recommended. Conformance to these guidelines will be recommended for national pilots. Please send comments to the authors or to the discussion group osi-ds@CS.UCL.AC.UK.

"OSI NSAP Address Format For Use In The Internet", R Colella, R Callon, 02/13/1991 <draft-ietf-osinsap-format-01.txt, or .ps>

The Internet is moving towards a multi-protocol environment that includes OSI. To support OSI, it is necessary to address network layer entities and network service users. The basic principles of OSI Network Layer addressing and Network Service Access Points (NSAPs) are defined in Addendum 2 to the OSI Network service definition. This internet draft recommends a structure for the Domain Specific Part of NSAP addresses for use in the Internet that is consistent with these principles.

"Guidelines for OSI NSAP Allocation in the Internet", Richard Colella, Ella Gardner, Ross Callon, , 04/25/1991 <draft-ietf-osinsap-internetalloc-01.txt, .ps>

The Internet is moving towards a multi-protocol environment that includes OSI. To support OSI in the Internet, an OSI lower layers infrastructure is required. This infrastructure comprises the connectionless network protocol (CLNP) and supporting routing protocols. Also required as part of this infrastructure are guidelines for network service access point (NSAP) address assignment. This paper provides guidelines for allocating NSAPs in the Internet.

"Building an Internet Directory using X.500", S. Kille, 01/07/1991 <draft-ietf-osix500-directories-01.txt, or .ps>

The IETF has established a Working Group on OSI Directory Services. A major component of the initial work of this group is to establish a technical framework for establishing a Directory Service on the Internet, making use of the X.500 protocols and services. This document summarises the strategy established by the Working Group, and describes a number of RFCs which will be written in order to establish the technical framework.

"The OSPF Specification, Version 2", John Moy, 01/23/1991 <draft-ietf-ospf-ospf2-01.txt, or .ps>

OSPF is a link-state based routing protocol. It is designed to be run internal to a single Autonomous System. Each OSPF router maintains an identical database describing the Autonomous System's topology. From this database, a routing table is calculated by constructing a shortest-path tree. OSPF recalculates routes quickly in the face of topological changes, utilizing a minimum of routing protocol traffic. OSPF provides support for equal-cost multipath. Separate routes can be calculated for each IP type of service. An area routing capability is provided, enabling an additional level of routing protection and a reduction in routing protocol traffic. In addition, all routing protocol exchanges are authenticated. This memo documents version 2 of the OSPF protocol. Version 1 was documented in RFC 1131. Distribution of this memo is unlimited.

"OSPF Version 2 Management Information Base", Rob Coltun, Fred Baker, 04/04/1991 <draft-ietf-ospf-ospfmib-03.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 defines objects for managing OSPF Version 2. Please send comments to ospf@trantor.umd.edu.

"Experience with the OSPF Protocol", Jon Moy, 04/01/1991 <draft-ietf-ospf-experience-00.txt, .ps>

This report documents experience with OSPF V2. This includes reports on interoperability testing, field experience, simulations and the current state of OSPF implementations. It also presents a summary of the OSPF Management Information Base (MIB), and a summary of OSPF authentication mechanism. Please send comments to ospf@trantor.umd.edu.

"OSPF Protocol Analysis", John Moy, 04/01/1991 <draft-ietf-ospf-analysis-00.txt, .ps>

This report attempts to summarize the key features of OSPF V2. It also attempts to analyze how the protocol will perform and scale in the Internet. Please send comments to ospf@trantor.umd.edu.

"Privacy Enhancement for Internet Electronic Mail: Part IV – Certifying Authority and Organizational Notary Services", Burt Kaliski, 08/14/1990 <draft-rsadsi-pemforms-01.txt>

This document describes two services that vendors may provide in support of Internet privacy-enhanced mail: certifying authority services on behalf of organizations, and organizational notary services for users. It also specifies the forms for interacting with vendors providing those services. This document is intended as a reference for vendors and for implementors of privacy-enhanced mail software; it is not at the appropriate level for users. The document also lists vendors.

"Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures", John Linn, 03/26/1991 <draft-irtf-psrg-pemmsgproc-00.txt>

This document defines message encryption and authentication procedures, in order to provide privacy-enhanced mail (PEM) services for electronic mail transfer in the Internet. Comments should be sent to <pem-dev@tis.com>.

"Definitions of Managed Objects for the Point-to-Point Protocol", Frank Kastenholz, 09/11/1990 <draft-ietf-pppext-pppmib-01.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 managed objects used for managing subnetworks using the Point-to-Point Protocol.

"The Point-to-Point Protocol Configuration Options: Negotiation of 32-bit FCS", Arthur Harvey, 12/20/1990 <draft-ietf-ppp-32bitconfig-01.txt>

This document defines a method to negotiate a 32-bit FCS Configuration Option for PPP. The Point-to-Point Protocol (PPP) provides a method for transmitting datagrams over serial point-to-point links. PPP is composed of three parts:

"The Point-to-Point Protocol: LLC over PPP", Arthur Harvey, 12/20/1990 <draft-ietf-ppp-lccoverppp-01.txt>

This document defines the operation of the LLC protocol over PPP. 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. The PPP encapsulating scheme, the basic LCP, and an NCP for controlling and establishing the Internet Protocol (IP) (called the IP Control Protocol, IPCP) are defined in RFC 1171 "The Point-to-Point Protocol for the Transmission of Multi-Protocol Datagrams Over Point-to-Point Links". IEEE 802.2 Logical Link Control (LLC) protocol provides additional services beyond those available directly from the various IEEE 802 Medium Access Control (MAC) data link protocols.

"Point-to-Point Protocol Extensions for DECnet Phase IV", Steven Senum, 06/04/1991 <draft-ietf-pppext-decnet-00.txt>

The purpose of this memo is to define a method for transmitting DNA Phase IV Routing packets over a serial link using the PPP protocol. This memo only applies to DNA Phase IV Routing messages (both data and control), and not to other DNA Phase IV protocols (MOP, LAT, etc). There are two basic approaches to running the DNA Phase IV Routing protocol over a serial line: 1. The approached that several router vendors have taken which is to treat the serial link as an Ethernet, using the same data and control messages an Ethernet would use. 2. The approach defined by Digital, which uses DDCMP and slightly different control messages. This memo will define a method that uses the first approach.

"ICMP Router Discovery Messages", S. Deering, 03/27/1991 <draft-ietf-rdisc-icmpmessage-00.txt>

This document specifies an extension of the Internet Control Message Protocol (ICMP) to enable hosts attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers. Please send comments to gw-discovery@gregorio.stanford.edu.

"Requirements for Internet IP Routers", Philip Almquist, 03/06/1991 <draft-ietf-rreq-iprouters-01.txt>

This draft attempts to define and discuss requirements for devices which perform the network layer forwarding function of the Internet protocol suite. The Internet community usually refers to such devices as "routers". This document is intended to provide guidance for vendors, implementors, and purchasers of IP routers.

"SNMP Over IPX", Raymond Wormley, 08/27/1990 <draft-ietf-snmp-snmpoveripx-00.txt>

The SNMP protocol has been specified as the official network management protocol of the Internet. Its widespread acceptance and implementation by developers, both inside and outside the Internet community, is fostering synergetic growth to a variety of protocols and platforms. This memo addresses the use of SNMP over Novell's proprietary IPX protocol. Roughly equivalent to UDP in function, IPX provides connectionless, unacknowledged datagram service over a variety of physical media and protocols.

"Definitions of Managed Objects for the Ethernet-like Interface Types", John Cook, 04/15/1991 <draft-ietf-snmp-ethernetmib-05.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 defines objects for managing ethernet-like objects.

"Use of the Community String for SNMP Proxys", Richard Fox, 12/31/1990 <draft-ietf-snmp-proxys-01.txt>

This memo defines an experimental portion of the Management Information Base (MIB) for use with network management protocols in TCP/IP-based internets.

"Definitions of Managed Objects for the SIP Interface Type", Kaj Tesink, 04/05/1991 <draft-ietf-snmp-smdsipmib-01.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 defines objects for managing SIP (SMDS Interface Protocol) objects.

"SNMP Communications Services", Frank Kastenholz, 04/23/1991 <draft-ietf-snmp-commservices-00.txt>

This Internet Draft is being distributed to members of the Internet community as an Informational RFC. The intent is to present a discussion on the issues relating to the communications services for SNMP. While the issues discussed may not be directly relevant to the research problems of the Internet, they may be interesting to a number of researchers and implementors.

"Comments on SNMP Proxy via Use of the @ sign in an SNMP Community", Jeff Case, et. al., 10/20/1990 <draft-ietf-snmp-proxycomments-00.txt>

This memo presents technical criticisms of introducing programmatically interpreted structure into the SNMP community string, as proposed in the Internet Draft entitled "Use of the Community String for SNMP Proxys".

"SNMP Administrative Model", James Galvin, 04/09/1991 <draft-ietf-snmpsec-admin-00.txt, .ps>

This paper presents an elaboration of the SNMP administrative model which provides a unified conceptual basis for administering SNMP protocol entities to support authentication and integrity, privacy, access control, and the cooperation of multiple protocol entities. This paper also describes how the elaborated administrative model is applied to realize effective network management in a variety of configurations and environments. The model described here entails the use of distinct identities for peers that exchange SNMP messages. Thus, it represents a departure from the community-based administrative model. By unambiguously identifying the source and intended recipient of each SNMP message, this new strategy improves upon the historical community scheme both by supporting a more convenient access control model and allowing for effective use of asymmetric (public key) security protocols in the future.

"Definitions of Managed Objects for Administration of SNMP Parties", Keith McCloghrie, James R. Davin, James M. Galvin, , 04/09/1991 <draft-ietf-snmpsec-mib-00.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 SNMP parties defined in "SNMP Administrative Model" as objects defined according to the Internet Standard SMI "Structure and Identification of management Information for TCP/IP-based internets" (RFC 1155). These definitions are consistent with the SNMP Security protocols set forth in "SNMP Security Protocol".

"SNMP Security Protocols", James M. Galvin, Keith McCloghrie, James R. Davin, , 04/09/1991 <draft-ietf-snmpsec-protocols-00.txt, .ps>

This draft document will be submitted to the RFC editor as a protocol specification. Distribution of this memo is unlimited. Please send comments to the authors: James M. Galvin galvin@tis.com, Keith McCloghrie kzm@hls.com, and James R. Davin jrd@ptt.lcs.mit.edu.

"Guidelines for the Secure Operation of the Internet", Richard Pethia, Steve Crocker, Barbara Fraser, , 04/01/1991 <draft-ietf-spwg-secureop-01.txt>

The purpose of this document is to provide a set of guidelines to aid in the secure operation of the Internet. Comments by Vinton G. Cerf, Vice President, Corporation for National Research Initiatives, and Chairman, Internet Activities Board, and James Van Bokkelen, President, FTP Software, Inc., have been provided to further illuminate the history and issues involved in this policy. Please send comments to spwg@nri.reston.va.us. This revision corrects a typographical error in the preamble of the document.

"Security Policy Handbook", Paul Holbrook, Joyce Reynolds, 05/31/1991 <draft-ietf-ssph-handbook-00.txt>

This handbook is the product of the Security Policy Handbook Working Group (SSPHWG), a combined effort of the Security Area and User Services Area of the Internet Engineering Task Force (IETF). This RFC provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

"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>

"Telnet Authentication Option", Dave Borman, 08/08/1990 <draft-ietf-telnet-authentication-01.txt>

"Telnet Environment Option", Dave Borman, 08/08/1990 <draft-ietf-telnet-environment-01.txt>

"FYI on an Internet Trouble Ticket Tracking System for addressing Internet User Connectivity Problems", M. Mathis, D. Long, 02/11/1991 <draft-ietf-ucp-connectivity-00.txt>

Users having trouble with the Internet are directed to contact their designated Network Service Center. The Network Service Center creates a Trouble Ticket which is registered with the Ticket Tracking System. The ticket is an agreement to obtain closure with the user. Network Service Centers can fix problems, track the work of others, or transfer responsibility for the ticket to other Network Service Centers using a formal hand-off procedure. Ticket hand-offs are coordinated by the Ticket Tracking System and ticket progress is monitored by the Ticket Support Centers. User complaints with the problem resolution process may be lodged with a Ticket Support Center, which will act on behalf of the user in resolving the problem.

"NOC Internal Integrated Trouble Ticket System Functional Specification Wish-list", Dale S. Johnson, 02/26/1991 <draft-ietf-ucp-tt-01.txt>

This Internet Draft describes general functions of a Trouble Ticket system that could be designed for Network Operations Centers. The document is being distributed to members of the Internet community in order to stimulate discussions of new production-oriented operator-level application tools for network operations. Hopefully, this will result both in more ideas for improving NOC performance, and in more available tools that incorporate those ideas. This memo does not specify a standard. Distribution of this memo is unlimited.

Chapter 2

Steering Group Report

2.1 Standards Progress Report

Between the December Meeting at the University of Colorado and the March meeting at Washington University in St. Louis, there have been many IETF originating protocols published as RFC's.

RFC1195	Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
	This RFC describes a proposed standard interdomain routing protocol for use in the Internet. It is the product of the IS-IS Routing Working Group.
RFC1206, FYI4	FYI on Questions and Answers - Answers to Commonly asked "New Internet User" Questions
	FYI 4 was updated and republished by the User Services Working Group.
RFC1207, FYI7	FYI on Questions and Answers - Answers to Commonly asked "Experienced Internet User" Questions
	This RFC is the product of the User Services Working Group.

2.2 Minutes of the Open Plenary and IESG

Agenda:

- The criterion for advancing routing protocols to Draft Standards
- The advancement of OSPF to draft standard state and,
- The advancement of BGP to the draft standard state.
- The beginning of work on IP version 7.

Criterion for Advancing Routing Protocols

Bob Hinden presented his criterion for the advancement of routing protocols to Draft Standard. The criterion enjoyed broad support in principle. The most interesting issues were:

- A clear definition of what it means for an implementation to be truly independent. It was felt that the code should have been written mostly from scratch yet it was agreed that no implementor can be expected to work in a complete vacuum.

- A definition of a "moderate" number of routers constituted was debated. There was a feeling that moderate should somehow simulate the size of a large corporate network. Because such a network cannot be built with older routing technology, no such networks exist for testing an implementation. Simulation was offered as a possible compromise. Many felt that actual experience at the expected size of production networks would be required for full standard status.

  In the end, after both points were debated, the Plenary accepted the time-honored principle of "I know it when I see it".

- The security requirements presented some interesting challenges. At this time, there are no definitive notions of what security means in terms of a routing protocol. The Plenary agreed that all features, including any defined security features, needed to be tested to reach draft standard.

Presentation to Advance OSPF to Draft Standard

John Moy, Jeff Burgan and Rob Coltun presented a detailed report on the deployment and testing experience of OSPF, as well as an exploration of the limits of the OSPF Routing protocol. (Slides are included later in these Proceedings)

After the presentation the Plenary discussed OSPF, and generally agreed that the protocol met the criterion for advancing to draft standard.

Presentation to Advance BGP to Draft Standard

Yackoff Rekter, Sue Hares and Dennis Ferguson presented a report on BGP, including operational experience and a discussion of the outer limits of the protocol. Several issues were raised which did not reach consensus.

- First was a reservation about the use of TCP as a transport protocol for BGP. Dennis Ferguson alluded to problems in conjection situations with TCP, where backoff may often result in degraded performance. Many felt that BGP should be re-written to use the UDP transport protocol. Others, especially those with ailing production networks felt a need to deploy BGP immediately as is. There were questions about whether BGP was a short-term solution or long-term, with many feeling the decision over whether to make BGP run over UDP depended on this analysis. The merits of these objections were debated, but not resolved.

- Second, the relationship between the IETF BGP protocol and the ANSI IDRP protocol was discussed. Many wanted assurances that the BGP protocol would remain stable and would not continue to incrementally evolve to conform to the emerging IDRP protocol. These concerns were addressed and generally satisfied the Plenary.

IP Version 7

A work item was proposed for a new Working Group to develop a new version of the IP specification. Goals of this revision were invisioned to include:

1. An expansion of the address space,
2. Some form of addressing structure to deal with very large networks, and
3. Resource allocation.

Noel Chiappa, Area Director for the Internet Services Area accepted an action item to form such a group.

Chapter 3

Area and Working Group Reports

3.1 Applications Area

Director(s):

- Russ Hobby: rdhobby@ucdavis.edu

Area Summary reported by Russ Hobby/UCDavis

Working Groups in Support of Multi-media

There were three Working Groups that met in St. Louis that will be closely coordinated to allow the use of multi-media in electronic mail on the Internet.

1. The Mail Extensions Working Group met and had lively discussions on how to update RFCs 821(SMTP) and 822 to allow the transfer of multiple "body parts" including binary data. The needed extensions will be designed to be as compatible with X.400 as possible. This should ease the implementation of RFC 821/822 to X.400 gateways.

2. A new Working Group in the OSI area was started to explore some testing of the Office Document Architecture (ODA) protocol on the Internet. ODA body parts will be defined by the Mail Extensions Working Group.

3. The FAX Working Group discussed the more general issue of what formats could be used for transmitting images across the Internet in electronic mail. Three possibilities are the FAX body part in X.400, images in the ODA format, and TIFF as specified by a document written at ISI. All efforts will be coordinated with the Working Groups above.

Other Working Groups Meeting in St. Louis

1. Domain Name System (dns) - This Working Group has a new mailing list since namedroppers was not specifically for the Working Group. The list is:

   dns-wg@nsl.dec.com for sending to the Working Group mail list

   dns-wg-request@nsl.dec.com to subscribe to the mail list.

   The main items covered by the Working Group were:
   - Discussion of Philip Almquist's project for a new BIND.
   - DNS MIB Variables.
   - Review DNS Security.
   - Collect information for the DNS Operators Guide.
   - DNS support for the Resource Location and Dynamic Hosts protocols.

2. Listserv BOF (listserv) - At the St. Louis meeting there was a Birds-of- a-Feather (BOF) session on the BITNET LISTSERV function to see how listserv works and explore some current needs in the Internet. David Lippke (UTexas) presented the

current state of listserv and answered many questions about it. It was determined that a Working Group should be created to define an Internet version of listserv.

3. Network Database (netdata) - This first meeting was to discuss the standardization of SQL databases operating over TCP/IP. Main points of work for the Working Group are to define a client/server model protocol over TCP/IP and provide security.

4. Network Printing Protocol (npp) - There were three main subjects discussed at this meeting.

   (a) Work was continued to document LPR/LPD as used today.
   (b) There was discussion of the Printer Access Protocol (PAP).
   (c) Discussion of how to do printer "spooling" and job submission.

5. Service Location Protocol (svrloc) - This was the first official meeting of this Working Group. The group decided that the first work item would be to define a protocol to locate resources, such as printers, connected to a site network. It was viewed that perhaps additions to the Domain Name System could solve some of this problem.

6. Telnet (telnet) - There was discussion of the Environment Option and advantages and disadvantages of user defined variables. Authentication and Encryption Options were the subject of the remainder of the meeting.

Working Groups Not Meeting in St. Louis

1. Distributed File System (dfs) - Although this group has not been too active and did not meet, new activity in this area has been found by another group, TSIG, in defining a trusted version of Network File System (NFS). The efforts of TSIG will be coordinated with the DFS Working Group.

2. Distributed Scheduling Protocol (chronos) - This Working Group did not meet in St. Louis but has made quite a bit of progress on the mail list in defining a protocol to allow scheduling between calendars maintained on different computers across a network. A draft document on the protocol has been written and is available as an Internet Draft.

CURRENT MEETING REPORT

Reported by David Lippke/UTEXAS

LISTSERV BOF Minutes

The LISTSERV BOF at the IETF meeting in St Louis went well. It was attended by a positive subset of the Internet Mail Extensions Working Group plus a number of people involved with the Applications Area of the IETF. In addition, Bruce Crabill, Larry Snodgrass, and John Wobus attended, therefore the full BITNET contingent was probably represented. However, it was clear that a number of the internetists had made an effort to educate themselves on LISTSERV and at least one had experience with maintaining a LISTSERV list.

David Lippke worked through about 35 foils which roughly followed the outline below:

1. BITNET background (size, growth, topology, traffic statistics).
2. LISTSERV background (history, present status, various statistics).
3. Definition of the original problem.
4. Overview of the main concepts, features, and facilities (13 points).
5. Two foils worth of list peering and its application.
6. Five foils on "distribute".
7. A bit on relayed file distribution.
8. Two foils on loop detection/suppression.
9. Six foils as a detailed look at list configuration.
10. A bit on NAD capabilities.
11. A slow three foils on user commands and options.
12. A bit on list archival and file storage (it was stressed several times throughout that standard list archival was a big deal).
13. Two foils on the database facility.
14. Then about three foils giving David's opinion of what's needed in the way of LISTSERV capabilities on the internet.

The whole session was rather interactive with at least a question or two asked on each foil and the responses to the questions occasionally turning into short discussions. It was a good crowd to talk to since everyone had already done some (or maybe a lot) of thinking about related issues. Consequently, there wasn't much left in regard to LISTSERV itself when David reached the end of his foils.

Phill Gross took over the discussion at that point and concluded with the group that a "LISTSERV" Working Group should be formed since there are clearly things which need to be done and there's also enough interest in the project to make things happen. One person said that they would like to see a basic set of user operations (upon lists) defined first so that they could proceed with a simple implementation of those while a compatible superset was developed over time. There seemed to be general agreement with this notion.

Phill then suggested that a list be formed to proceed with the Working Group discussion before the July IETF. Chris Myers of Washington University volunteered listdev@wugate.wustl.edu and this was accepted. To subscribe send e-mail to listserv@wugate.wustl.edu with a body part of "add listdev". To unsubscribe send e-mail to listserv@wugate.wustl.edu with a body part of "delete listdev". To send a message to the list send e-mail to listdev@wugate.wustl.edu. The initial membership will be taken from the attendance roster with the addition of Eric. Anyone who is interested in participating in the Working Group is welcome to join the list

A PostScript copy of the foils is available upon request from David Lippke.

Attendees

Bruce Crabill	bruce@umdd.umd.edu
David Crocker	dcrocker@pa.dec.c
Ralph Droms	droms@bucknell.edu
Johnny Eriksson	bygg@sunet.se
Demi Getschko	"DEMI@FPSP.HEPNET"
Phillip Gross	pgross@nri.reston.va.us
Russ Hobby	rdhobby@ucdavis.edu
Neil Katin	katin@eng.sun.com
Shelly Knueven	shelly@wugate.wustl.edu
Vincent Lau	vlau@sun.com
Eliot Lear	lear@net.bio.net
David Lippke	lippke@utdallas.edu
Chris Myers	chris@wugate.wustl.edu
Lawrence Snodgrass	snodgras@educom.edu
Bernhard Stockman	bygg@sunet.se
Gregory Vaudreuil	gvaudre@nri.reston.va.us
John Wobus	jmwobus@suvm.acs.syr.edu

A Technical Introduction to LISTSERV

David Lippke
The University of Texas at Dallas
lippke@utdallas.edu
+1 214-690-2632

Background info on BITNET

- Nearly 3400 nodes on the network
  - About 1000 IBMs running VM
  - About 1700 Vaxen
  - Rest are IBM/MVS, Unix, CDC, HP, Honeywell, etc...
- Linear growth for last several years
- Topology —
  - 7.15 hops between nodes on average
  - 17 hop network diameter
  - Approx 1.12 * N links

BITNET Background, cont'd

- 75% of files are less than 5K bytes; 25% less than 2K bytes
- Major sites push 300 to 1000 megabytes per day; 75K to 250K files per day
- These sites are generally near NSSes with interconnections following NSFnet topology

LISTSERV Background

- Written by Eric Thomas, then a student at Ecole Centrale de Paris, now with SUNET and Nordunet
- Written in REXX and assembler
- First deployed during the summer of 1986
- Very fast paced development through 1987
- Currently updated about every six months
- Main discussion list is LSTSRV-L at POLYGRAF, RUTVM1, SEARN, and UGA

LISTSERV Background, cont'd

- 257 registered LISTSERVs
- 137 "backbone" LISTSERVs
- 22 list database LISTSERVs
- Approximately 2000 public lists
- Major sites see 2000–2500 jobs per day
- Secondary sites see 700–1500 per day

Original Problems Addressed

- Usability problems with the primitive "LISTSERVs" of the day — similar to the current situation with internet-based lists
- Very simplistic mail explosion and network congestion as a result
  - No distribution of the work
  - Individual copies of each posting sent to each list member, even for members on the same node.
  - Very similar in effect to current internet list handling

Session Outline

- ☐ Concepts, features, and facilities
- ☐ Summary of LISTSERV facilities which may be needed on the internet
- ☐ Additional things which could/should be done and things which should avoided
- ☐ Discussion

Concepts, Features, and Facilities

- ☐ Mail "from" the list ID
- ☐ Peered lists
- ☐ Mail via "distribute"
- ☐ Relayed file distribution
- ☐ Loop detection and traffic control
- ☐ Very flexible list control and handling
- ☐ Facilities for system administrators
- ☐ User facilities

Concepts/features/facilities, cont'd

- ☐ File and list archive storage
- ☐ Database facilities
- ☐ Optional automatic code maintenance
- ☐ Automatic database updating & maint
- ☐ User Database

Peered Lists

- ☐ Splits lists over two or more nodes
- ☐ Users are subscribed to the closest node
- ☐ Submissions can be sent to any peer
- ☐ Many administrative facilities, but still somewhat difficult to manage
- ☐ Application —
  - Huge lists
  - Multiple archive sites
  - Good for getting postings over links with high backlogs or bad performance

Sample Peered List — RSCSMODS

```
                     EBOUB011 <--> FINHUTC
                                      |
                                      V
TAMVM1 <--> OHSTVMA <--> UBVM <--> POLYGRAF
               |
               V
              UGA
```

Mail via "Distribute"

- ☐ Distributes list traffic on the basis of minimizing total network load.
- ☐ The initial LISTSERV decides how to distribute the work.
- ☐ "distribute" jobs are either sent directly to the final backbone nodes or are grouped together and sent to intermediate nodes.
- ☐ Most small to medium (50–750 member) lists are handled this way.

Relayed File Distribution

- ☐ Any user can submit distribution jobs to the backbone LISTSERV network.
- ☐ Employs a special JCL which is used to specify the recepients and processing options; the file to be distributed is appended to this.
- ☐ Typically used by software maintainers to distribute updates or to distribute large files like newsletters.

Loop Detection

- ☐ Examines the subject field for known trouble subjects.
- ☐ Examines the from field for known bad userids like 'mailer', '*daemon*', and 'postmaster'
- ☐ Looks inside the body for a copy a message with fields referencing the list userid.
- ☐ Does NOT use 'Message-ID:' — runs a CRC on the stripped body and throws away anything with the same signature as one of the previous N messages.

Loop suppression, Traffic control

- ☐ Per-list daily limits on message traffic
- ☐ Constant sources of bad commands are "served off" and silently ignored until the postmaster restores service.
- ☐ Constantly monitors link backlogs — goes offline and online like Unix acct(2).

List Configuration Terms

- ☐ (access-level)
  - Public
  - Postmaster
  - A1,A2,... where Ai is
    1. Private
    2. (listname)
    3. Owner or Owner(listname)
    4. Service or Service(listname)

List Configuration Terms

- ☐ (destination)
  - List
  - Sender
  - Both
  - None
  - "address"
- ☐ (interval)
  - Yearly, Monthly, Weekly, Daily, Hourly
  - Single

List Configuration Terms

- ☐ (area)
  - Name of a network
  - Name of a country
  - Value of the 'local' keyword
  - A nodename or pattern
- ☐ (mon-address)
  - "address"
  - 'Postmaster' or 'Postmasters'
  - 'Owner' or 'Owners'

Major List Keywords

- ☐ Review= (access-level)
- ☐ Subscription= By_owner | Open | Closed
- ☐ Send= (access-level) | Editor
- ☐ Reply-to= (destination),Respect | Ignore
- ☐ Files= Yes | No
- ☐ Confidential= No | Yes | Service
- ☐ Validate= Store only | All commands
- ☐ X-Tags= Yes | No | Comment

Major List Keywords

- ☐ Stats= Normal | Extended | None, (access-level)
- ☐ Ack= Yes | Msg | No
- ☐ Notebook= No | (Yes, (fm), (interval) | Separate, (access-level))
- ☐ Owner= (address1) | (access-level1), ...
- ☐ Editor= (address1), ...
- ☐ Language= idiom

Major List Keywords

- ☐ Service= (area1), ...
- ☐ Local= node1, ...
- ☐ Errors-To= (mon-address1), ...
- ☐ Default-Options= user_options

Node Administrator Facilities

Node Administrators can —

- Submit commands on the behalf of any local user.
- Can have a users automatically deleted from all lists they are subscribed to (in the world).
- Hold / Free lists
- Put the LISTSERV offline, etc.

User List Options

Controlled by each user on a per list basis —

- Mail / NoMail
- Files / NoFiles
- Ack / NoAck / MsgAck
- Repro / NoRepro
- Conceal / NoConceal

Major User Commands

- ☐ Help / Info ?
- ☐ List [detailed | long | short]
- ☐ Query listname
- ☐ SUBscribe/SIGNON listname [fullname]
- ☐ UNSubscribe/SIGNOFF listname
- ☐ SET listname options
- ☐ REView listname [(options]

Major User Commands

- ☐ STats listname [(local]
- ☐ GET filename
- ☐ PUT filename [options]
- ☐ INDex filelist_name
- ☐ PW add/change/delete password(s)
- ☐ AFD/FUI add/delete/get
- ☐ SHOW [options]

List Archival / File Storage

- ☐ Lists are achived and available according to their definition. Users can fetch entire logs or have entries sent to them which have been flagged by database operations.
- ☐ Files are organized under "filelists" and either referenced explicitly or as a part of "packages."
- ☐ Users can subscribe to packages or individual files.
- ☐ GET and PUT permissions are controlled at all levels by access lists.

Database Facility

- ☐ General database engine used to query the master BITNET node database, LISTSERV node database, and any list archives
- ☐ Access primarily through interactive messages and frontend programs which send the raw commands and parse the results.
- ☐ Complex English queries allowed
- ☐ Most often used to lookup information in a list archive and have the "hits" sent back to the user.

Database Query Examples

```
Search Rosemary in MOVIES

Search Hardware problem with a 4381 in IBM-L

Search wooden chair (blue or green) in CHAIRS

Search problem in BBOARD since July

Search place chair in CHAIRS where price < 50

Select * in BITEARN where -
            site sounds like COHRNEAL

Select * in TCP-IP where -
            subject contains fax or facsimile
```

What's needed on the Internet?

Parts of most everything, although —

- ☐ Peering is probably best left to die. Netnews handles the large list distribution.
- ☐ File distribution is difficult, but there are perhaps some creative options available.
- ☐ The "distribute" function is also difficult, but very important.
- ☐ User directory service is best treated as an independent resource of which the LISTSERV function would be a client.

Thoughts on Design & Additional Functions

- ☐ Slicing, dicing, megaserver should probably be avoided
- ☐ The functions and specifications should be modularized and split out, but it'd be nice if they were able to use each other to present a well-integrated overall facility.
- ☐ ???

3.1.1 Distributed File Systems (dfs)

Charter

Chair(s):
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 Working Group 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 Working Group 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.

3.1.2 Distributed Scheduling Protocol (chronos)

Charter

Chair(s):
Paul Linder, `lindner@boombox.micro.umn.edu`

Mailing Lists:
General Discussion: `chronos@boombox.micro.umn.edu`
To Subscribe: `chronos-request@boombox.micro.umn.edu`

Description of Working Group:

The Chronos protocol Working Group is chartered to define a protocol for the management of calendars, appointments and schedules over the internet. In defining this protocol, several questions must be addressed. The role of the calendar administrator must be defined. Differing levels of security need to be specified to allow maximum functionality yet still allow privacy and flexibility. The scope of the protocol should also be evaluated; how much burden should we put on the server, on the client? Additionally the behavior of multiple chronos servers must be analyzed.

This protocol should be able to be developed and stabilized within 6-8 months, since there is already a draft specification to work from. The process is subject to extension if many new features are added, or more revision is needed.

Goals and Milestones:

Jan 1991 Review first draft document, determine necessary revisions. Follow up discussion will occur on mailing list. Prototype implementations.

Feb 1991 Make document an Internet Draft. Continue revisions based on comments received over e-mail.

Mar 1991 Spring IETF meeting. Review final draft and if OK, give to IESG for publication as RFC. Begin implementations.

Jul 1991 Revise document based on implementations. Ask IESG to make the revision a Draft Standard.

3.1.3 Domain Name System (dns)

Charter

Chair(s):
Michael Reilly, reilly@pa.dec.com

Mailing Lists:
General Discussion: dns-wg@nsl.dec.com
To Subscribe: dns-wg-request@nsl.dec.com

Description of Working Group:

The DNS Working Group is concerned with the operation of name servers on the Internet. We do not operate name servers but serve as a focal point for the people who do operate them. We are also concerned with the Domain Name System itself. Changes to the existing RFC's, for example, are discussed by the Working Group. If changes to the RFC's or additional DNS related RFC's are deemed necessary the Working Group will propose them and will prepare the associated documents.

Because we intend to serve as the focal point for people operating name servers, one of our projects will be to assist anyone bringing up a name server by publishing a collection of useful hints, tips and operational experience learned by the people already running name servers.

The DNS Working Group will also take an active role in the dissemination of solutions to problems and bugs encountered while running various name server implementations. We will also provide guidance to anyone writing a new name server implementation, whenever possible.

Goals and Milestones:

TBD	Adding DNS variables to the MIB.
TBD	Hints, tips, and operations guide for DNS software
TBD	Implementation catalog for DNS software.
TBD	Discussion of adding load balancing capability to the DNS.
TBD	Discussion of adding a Responsible Person Record.
TBD	Discussion of adding network naming capability to the DNS.
TBD	Evaluate short-term measures to improve, or at least describe the security of the DNS.

CURRENT MEETING REPORT

Reported by Michael Reilly/DEC

DNS Minutes - Boulder

This was sort of a startup meeting for the Domain Name System (DNS) Working Group. Our primary goal was to determine the short and long-term goals of the group and to determine what direction the group wants to take.

We began by discussing a list of potential areas for the Working Group to focus on. We also discussed some of the problems in the DNS and the BIND nameserver implementation. Several areas were identified for further work.

In the near term the group will update the Charter and discuss the group's goals on the mailing list.

In the immediate future we will concentrate on:

1. Increasing the robustness of the existing BIND implementation by fixing a couple of known problems and working with Berkeley (UCB) to encourage use of the updated BIND code.
2. Increasing the robustness of the DNS by encouraging root nameservers to incorporate the updated BIND code.
3. Increasing the robustness of the DNS by encouraging root nameservers to NOT preform recursive lookups.
4. Collecting as much of the existing wisdom concerning management and operation of a DNS server as well as BIND specific information as we are able to. It is anticipated this information will be made available in the for of a "cookbook" style document.

Attendees

Steve Alexander	stevea@i88.isc.com
Philip Almquist	almquist@jessica.stanford.edu
William Barns	barns@gateway.mitre.org
Robert Collet	/pnrobert.d.collet/o=us.sprint/admd=telemail/c=us/@sprint.com=
Curtis Cox	zk0001@nhis.navy.mil
Vince Fuller	vaf@Standford.EDU
Robert Gilligan	gilligan@sun.com
Juha Heinanen	jh@funet.fi
Darren Kinley	kinley@crim.ca

Holly Knight	holly@apple.com
Alex Koifman	akoifman@bbn.com
E. Paul Love	loveep@sdsc.edu
Paul Mockapetris	pvm@darpa.mil
Lynn Monsanto	monsanto@eng.sun.com
Michael Reilly	reilly@pa.dec.com
Robert Reschly	reschly@brl.mil
Tim Seaver	tas@mcnc.org
Richard Smith	smiddy@pluto.dss.com
Roxanne Streeter	streeter@nsipo.nasa.gov
Sally Tarquinio	sally@gateway@mitre.org
Glenn Trewitt	trewitt@pa.dec.com
A. Lee Wade	wade@discovery.arc.nasa.gov
Walter Wimer	walter.wimer@andrew.cmu.edu
Cathy Wittbrodt	cjw@nersc.gov

CURRENT MEETING REPORT

Reported by Michael Reilly/DEC

DNS Minutes

We began this meeting of the DNS Working Group by discussing administrative details. Several Working Group members expressed problems in getting added to the namedroppers mailing list. The Working Group decided it would be beneficial to setup a DNS Working Group mailing list. Mike Reilly volunteered to host the list.

The mailing list is dns-wg@nsl.dec.com. Add/drop requests are to be sent to dns-wg-request@nsl.dec.com.

We spent some time discussing the Charter of the Working Group. Several goals were identified for the near term. Discussion of these goals is to take place on the mailing list. The top four goals are:

- Define DNS related MIB variables. The Working Group does not have experience in writing MIB definitions so we will produce a list of variables and work with others to produce the document describing them.

- Produce a DNS Operator's Guide for publication as an informational RFC. There was some discussion concerning the inclusion of BIND specific information in the RFC or whether the guide should simply discuss operational experience from the Internet. This was not resolved at the meeting but will be resolved as work progresses. The relationship between this document and RFC 1033 (Domain Operations Guide) as also discussed.

- Investigate additional Resource Records as well as other changes to the DNS RFC's. Many of the proposed changes to the DNS RFC's come as a of the publication of the host requirements RFC (1123) or as a result of operational requirements in the Internet. The Working Group expects to be involved in producing additional RFC's as the need arises.

- Investigate and document the security (or lack of) in the current DNS and in common implementations (for example, BIND).

Mike Reilly is working on an updated Charter which will describe these goals and include a time schedule. This will be posted to the mailing list for discussion and review next week.

The DNS Working Group will also be investigating ways to work closely with the individuals running DNS nameservers on the Internet. We hope to both learn form them so that the Operations Guide mentioned above will be as useful as possible and to serve as a useful resource for identification and resolution of problems encountered in operating DNS nameservers.

Phil Almquist is being funded to produce a more robust version of BIND for general distribution. He presented his plans and time schedule in detail during the DNS Working Group meeting. The discussion following his presentation indicated that there was positive support for his efforts.

We spent that last portion of the meeting discussing ways in which the DNS could and should be used as a part of dynamic host configuration. Several members of the Dynamic Host Configuration Working Group described their needs and limitations in the current DNS which prevent its use for this purpose. The discussions helped make many of the DNS Working Group members aware of the needs hosts wishing to dynamically configure. We will work with the Dynamic Host Configuration Working Group to determine how the DNS can help them.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Lida Carrier	lida@apple.com
Bill Durham	durham@MDC.COM
Robert Enger	enger@seka.scc.com
Charles Fineberg	fineberg@wums2.wustl.edu
Karen Frisa	karen.frisa@andrew.cmu.edu
Jeremy Greene	greene@coral.com
Patrick Heisinger	heisinger@mdcguy.mdc.com
Ajay Kachrani	kachrani@regent.enet.dec.com
Tom Kessler	kessler@sun.com
Anders Klemets	klemets@cs.cmu.edu
Bill Manning	bmanning@houston.sc.ti.com
Mike Marcinkevicz	mdm@calstate.edu
Carol Melowitz	melowitz@mdcgwy.mdc.com
Donald Merritt	don@brl.mil
Mark Moody	ccmarkm@umcvmb.missouri.edu
Michael Patton	map@lcs.mit.edu
Jan Michael Rynning	jmr@nada.kth.se
Tim Seaver	tas@mcnc.org
John Veizades	veizades@apple.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov
Wing Fai Wong	wfwong@malta.sbi.com

3.1.4 Internet Mail Extensions (smtpext)

Charter

Chair(s):
Gregory Vaudreuil, gvaudre@nri.reston.va.us

Mailing Lists:
General Discussion: ietf-smtp@dimacs.rutgers.edu
To Subscribe: ietf-smtp-request@dimacs.rutgers.edu

Description of Working Group:

The SMTP extensions Working Group is chartered to develop extensions to the base SMTP protocol (RFC821) to facilitate the more efficient transmission of 8 bit text and binary data. Among the extensions to be considered to SMTP are the elimination of the ASCII text character restriction and line length restriction to allow the sending of arbitrary 8 bit character sets, and the definition of mechanisms to facilitate binary transmission, and extensions to the negotiation sequence to facilitate batch transmission.

Goals and Milestones:

Apr 1991 Review the Charter of the group. Determine if changes to SMTP are necessary. Discuss the needs for backward compatability, and interoperabiliy. This discussion will be held by email.

Aug 1991 Discuss the elimination of the 7 bit restrictions in SMTP, and the implications of removing this restriction in terms of interoperation.

Aug 1991 Discuss the issues involved with binary transmission. Determine whether a "binary" mode should be pursued, and whether the SMTP line length restriction should be eliminated.

Dec 1991 Write a document specifying the changes to SMTP agreed to by the group. Post as an Internet Draft.

Mar 1992 Review and finalize the SMTP Extentions document.

Mar 1991 Submit the SMTP Extentions document as a proposed standard.

CURRENT MEETING REPORT

Reported by Greg Vaudreuil/CNRI

SMTPEXT Minutes

Agenda

- Introduction
- Why Are We Here?
- Should We Be Here?
- Goals For The Group
- Mail Extensions Architecture
- Message Format Architecture

The IETF Internet Mail Extensions Working Group met for two days at the 20th IETF meeting in St. Louis.

The meeting began with an overview of the motivations for forming the Working Group, and a discussion of the role the group should play in the context of the current Internet mail environment and the emergence of X.400 based mail systems. There was little debate about the necessity to engineer a short-term solution to the need for greater mail functionality, especially for international character set support. There was a feeling that the work of this group could potentially speed the X.400 deployment into the current Internet. By increasing the functionality of X.400 gateways and stimulating the development of multi-media mail facilities, the work may facilitate the smooth transition to X.400. No one expressed an opinion that this work should not continue.

The Working Group spent the remainder of the morning enumerating possible goals for the mail extensions effort. The group proceeded to narrow the list of goals to a manageable subset for the first phase of the effort.

Possible Goals

Goals chosen for the initial effort marked with an X.

x	Include support for most international multi-character sets in message body.
x	Support multi-part messages.
x	Support multi-media messages.
x	Increase interworkability with X.400.
x	Remain backward compatible with RFC 822, 821.
x	Support enhanced functionality over current 7 bit transport.
?	Use 8 bit transport paths if available.
?	Enhance multi-character set support in message headers.
x	Resolve line length, end of message, and format effector issues.
-	Resolve message length issues (Message Fragmentation).
-	Include external references for long messages.

- Define standard error message reporting formats (Internet Mail Control Message Protocol).
- Define a standard User Agent (UA) configuration file format (.mailcap).
- Mail Gateway requirements document.
- Receiver initiated file transfer.
- POP-IMAP-PCMAIL standardization issues.
- Subsume X.400 Functionality (Return Receipt, Privacy Enhanced Mail, Accounting).
- Listservice Specification.

? Mail Transport MIB.
? Enhanced addressing (i.e., Phone Number, Postal Address).
- Mailbox Management.
- Message Storage Architecture.

x Establish Liaison with X.400.

After enumerating the goals for the mail extensions effort, the group proceeded to categorize the goals as either RFC 822 Message Format Extensions or RFC 821 SMTP Extensions. The group briefly discussed the differences between RFC 821 and RFC 822, resulting in greater understanding of the current mail environment. One crucial distinction was the point in the specifications where ASCII-7 is defined to be the character set. It was found that SMTP does indeed specify ASCII as the character set, rather than the set of allowed bit codes.

Architecture

The Working Group proceeded to spend the second full afternoon session discussing the transport architecture to be used in enhancing the current Mail system. The architecture discussion was crucial to understand the context of the changes needed to the message format, and SMTP RFC's. Initially there were two competing ideas for this architecture, and later, a transition solution was proposed.

The 7 Bit Solution

The first proposal, based on the existing 7 bit infrastructure, specified no changes to the SMTP protocol, and made all mail functionality enhancements in the RFC 822 message format. In the special case of 8 bit text, the conversion to a 7 bit encoding occurs in the sending and receiving User Agents.

```
+---------+------+                              +------+----------+
| 8 Bit UA | 8to7 |                              | 7to8 | 8 Bit UA |
+---------+------+                              +------+----------+
                |                                      |
                |   +------------+     +-----------+   |
                +---| 7 Bit MTA  |-----| 7 Bit MTA |---+
                    +------------+     +-----------+
```

The 8 Bit Solution

The second proposal, based on current practice among those currently using extended character sets in Europe, consisted of lifting the 7 bit restriction in SMTP, and using existing 8 bit friendly User Agents to pass 8 bit character codes to capable terminals. This proposal has been referred to as the "declare 7 bit to be broken". It was asserted that most SMTP Message Transfer Agent's (MTA) currently pass 8 bit mail unmodified. This proposal requires no special encoding of 8 bit text.

```
+----------+    +------------+    +-----------+    +----------+
| 8 Bit UA |----| 8 Bit MTA  |----| 8 Bit MTA |----| 8 Bit UA |
+----------+    +------------+    +-----------+    +----------+
```

These two proposals are not interoperable. The first, the 7 bit solution, interoperates with current SMTP agents, but not with existing 8 bit users or their agents. The second works with existing 8 bit User Agents but not fully conformant SMTP implementations.

The 8/7 Bit Transition Solution

After some discussion, a transition solution was proposed by the Chair, soon to be dubbed the "Wretched" solution. This proposal required 8-bit capable SMTP agents to convert from 8 bit to 7 bit message formats. This proposal was based on the principle that a conversion from 8 bits to 7 bits can be specified such that the same conversion can be made either by a User Agent, or by a mail forwarder on a per-message demand basis.

This transition proposal has two distinguishing features. In the existing world of 7 bit SMTP MTA agents, it is identical to the 7 bit proposal, requiring all UA's to either encode or decode 8 bit text. In the ideal world where all SMTP MTA's are 8 bit capable, it is identical to the 8 bit solution. It does however require implementing the conversion process in both the MTA's and UA's.

A third feature, one that turned out to cause problems, is the requirement that the entire message be convertible from 8 bit to 7 bit without regard to the contents. It was felt that if a suitable encoding was chosen, it could be indicated by prepending a new header line "Message encoded in 7 bits" by any MTA that modified the message.

```
                 +------------+         +-----------+
    +--------->-| 8 Bit MTA  |--------| 8 Bit MTA | ------------+
    |            +------------+         +-----------+            |
    |                                                            |
    |      +------------+------+       +-----------+             |
    |    +-| 8 Bit MTA  | 8to7 |-------| 7 Bit MTA |---+         |
    |    | +------------+------+       +-----------+   |         |
    |    |                                             |         |
+----------+------+                                +------+----------+
| 8 Bit UA | 8to7 |                                | 7to8 | 8 Bit UA |
+----------+------+                                +------+----------+
           |                                           |
           |   +------------+       +-----------+      |
           +---| 7 Bit MTA  |-------| 7 Bit MTA |--+
               +------------+       +-----------+
```

At the conclusion of the first day, the group tentatively adopted the transition solution.

Day 2

The second day was scheduled to begin work on the specifics of the Message Format Extensions required to achieve the goals previously defined. The work was intended to be essentially independent of the RFC 821 SMTP efforts to be discussed later in the day. However, within minutes, it became clear that the group had not realized many of the implications of the transition proposal. Specifically, there is an implication that non-text messages originating from an 8 bit User Agent may, with certain encodings, be re-encoded by the MTA, resulting in double-encoding. For a worst case example, consider a binary message encoded to utilize a full 8 bit path. If it encounters a 7 bit MTA later in the journey, it will be converted again. While judicious choice of encodings will make this double encoding a non-issue, the perceived additional complexity, and the restrictions this implied in the multi-part, multi-media extensions to be proposed caused many in the group to re-evaluate their positions with regard to the transition proposal.

For the purpose of making progress the Working Group adopted the 7 bit proposal to begin work on the 822 message body extensions. There remains significant constituency for the transition proposal, but after hours of hallway discussions, the group reached a consensus that changes to SMTP merely to facilitate the 8 to 7 conversion were not sufficient to justify upgrading the MTA infrastructure. However, many hold hope that enhancements including binary transmission will result in a system that can fully and efficiently utilize 8 bit transport.

Message Format Extensions

After the contentious issues of mail transport were put behind the group, work began on defining an extension to the RFC 822 message format to facilitate multi-part, multi-media applications, including international character sets. The group began by considering a specific proposal by Borenstein, Freed, Vance, and Carosso (BFVC). As this proposal was put forth, a debate ensued over the relative merits of line counts vs message boundary delimiters. The group felt that in general, message delimiters were superior to line counts for reliability and readability, but that line counts were useful "hints" which allowed fast parsing of long multi-part messages. A proposal to combine both message delimiters and line counts was made, but not pursued.

The group moved forward and chose to use the BFVC proposal as a strawman. Several issues were raised.

The message boundary delimiter is chosen at random for each message. This eliminates the need to reserve a specific begin and end sequence for messages. It was not clear how difficult it would be to implement this scheme.

The content-encoding and content-type are independent fields which are included for each of the message body parts. Advocates asserted that these independent axis make the overall implementation easier than defining a standard encoding for each body part. This proposal allows a sender to encode a message in whatever encoding type is optimal for the message sent. The receiver must then be able to decode each of the several standard encoding types. With several standard encoding types defined, a sender could pick the ideal encoding for the particular message type. This many-types, limited encodings approach reduces the complexity for a full featured User Agent. This proposal has the disadvantage of increasing complexity in a single function station, such as a fax server, or text only User Agent.

The implication that a User Agent must implement several decoding and encoding mechanisms to simply receive and send 8 bit text was of some concern. This was discussed but not resolved. One proposal was to make 8 bit text a special case with a single encoding type.

A strawman poll was taken with the following options.

1. Body part "a" must be sent with encoding type "y".
2. Body part "a" should be sent with encoding type "y", but may be sent with any encoding x,y,z.
3. Body part "a" can be sent with any encoding x,y,z.
4. Body parts a, b, c can be sent in any encoding x,y,z except for body part "d" which must be sent in "x".

There was no majority, with most expressing preference for (2), and and equal number expressing either (3) or (4).

Future Meetings

The Chair of the Working Group strongly advocated an interim meeting. He proposed a choice between a face to face meeting or a Video Teleconference. The group preferred a Video Teleconference. The Chair took an action to find open dates and if possible, schedule a Video Teleconference. Interest was expressed by some of the international participants in holding a Working Group meeting in Europe in the near future.

Attendees

Nathaniel Borenstein	nsb@thumper.bellcore.com
Cyrus Chow	cchow@orion.arc.nasa.gov
Alan Clegg	abc@concert.net
Mark Crispin	mrc@cac.washington.edu
David Crocker	dcrocker@pa.dec.c
Johnny Eriksson	bygg@sunet.se
Ned Freed	net@ymir.claremont.edu
Shari Galitzer	shari@gateway.mitre.org
Tony Genovese	
Olafur Gudmundsson	ogud@cs.umd.edu
Russ Hobby	rdhobby@ucdavis.edu
Neil Katin	katin@eng.sun.com
Tom Kessler	kessler@sun.com
Darren Kinley	kinley@crim.ca
Anders Klemets	klemets@cs.cmu.edu
Jim Knowles	jknowles@trident.arc.nasa.gov
Ruth Lang	rlang@nisc.sri.com
Vincent Lau	vlau@sun.com
David Lippke	lippke@utdallas.edu
E. Paul Love	loveep@sdsc.edu
Leo McLaughlin	ljm@ftp.com
David Miller	dtm@ulana.mitre.org
Mark Moody	ccmarkm@umcvmb.missouri.edu
Keith Moore	moore@cs.utk.edu
Robert Morgan	morgan@jessica.stanford.edu
Michael Patton	map@lcs.mit.edu
Joel Replogle	replogle@ncsa.uiuc.edu
Jan Michael Rynning	jmr@nada.kth.se
George Sanderson	sanderson@mdc.com
Ursula Sinkewicz	sinkewic@decvax.dec.com
Lawrence Snodgrass	snodgras@educom.edu
Bernhard Stockman	bygg@sunet.se
Dean Throop	throop@dg-rtp.dg.com
Robert Ullmann	ariel@relay.prime.com
Gregory Vaudreuil	gvaudre@nri.reston.va.us
John Veizades	veizades@apple.com
Chris Weider	clw@merit.edu
John Wobus	jmwobus@suvm.acs.syr.edu
Russ Wright	wright@lbl.gov
Wengyik Yeong	yeongw@psi.com

3.1.5 Internet Message Extentions (822ext)

Charter

Chair(s):
Gregory Vaudreuil, gvaudre@nri.reston.va.us

Mailing Lists:
General Discussion: ietf-822@dimacs.rutgers.edu
To Subscribe: ietf-822-request@dimacs.rutgers.edu

Description of Working Group:

This Working Group is chartered to extend the RFC 822 Message format to facilitate multi-media mail and alternate character sets. The group is expected to formulate a standard message format, roughly based on either RFC1154 or RFC 1049. The immediate goals of this group are to define a mechanism for the standard interchange and interoperation of international character sets.

Goals and Milestones:

Done	Review the Charter, and refine the groups focus. Decide whether this is a worthwhile effort.
Done	Discuss, debate, and choose a framework for the solution. Assign writing assignments, and identify issues to be resolved.
Jul 1991	Review exiting writing, resolve outstanding issues, identify new work, and work toward a complete document.
Nov 1991	Post a first Internet Draft.
Dec 1991	Review and finalize the draft document.
Jan 1992	Submit the document as a Proposed Standard.

3.1.6 Network Database (netdata)

Charter

Chair(s):
Daisy Shen, daisy@watson.ibm.com

Mailing Lists:
General Discussion: ietf-ndb@ucdavis.edu
To Subscribe: ietf-ndb-request@ucdavis.edu

Description of Working Group:

The Network Database Working Group is chartered to define a standard interface among databases on TCP/IP networks. The Working Group will address the issue of database connectivity in a distributed environment which allows authorized users remote access to databases. It will be designed as a client/server model based on TCP/IP as its communication protocol.

Several problems must be resolved that are associated with the network database protocol, such as management of multiple threads between clients and servers, management of multiple servers, management of data buffers, data conversions, and security.

Additional related problems will be covered as the discussion goes on. Therefore, the description and the schedule can be revised.

This Working Group is independent from the SQL access group; however, there may be some overlapping interest. The SQL access group is welcome to join IETF's discussions and share information in both directions. If both groups find that merging two efforts in one will speed up the process, the merge can be done in the future. For now, this Working Group works on issues according to our own schedule and efforts.

Goals and Milestones:

Done — Review and approve the Charter, making any changes necessary. Examine needs, resources for this network database protocol and define the scope of work. Begin work on a framework for the solution. Assign writing assignments for first draft of the document.

Jun 1991 — First draft to be completed.

Aug 1991 — Review first draft document, determine necessary revisions. Discuss problems remained unsolved from the first IETF meeting.

Dec 1991 — Continue revisions based on comments received at meeting and e-mail. Start making document an Internet Draft.

Mar 1992 Review final draft. If it is OK, give it to IESG for publication as RFC.

Jun 1992 Revise document based on implementations. Ask IESG to make the revision a Draft Standard.

CURRENT MEETING REPORT

Reported by Daisy Shen/IBM

Network Database Minutes

This is a new Working Group Chaired by Daisy Shen. The first meeting consisted of discussing the Charter, defining the scope and preparing for the next meeting.

All attendees agreed upon the Charter; although there was no discussion regarding the schedule and the milestones. Those items will be discussed at the next meeting.

Define the Scope

1. It will be built as a client/server model which will be called a database requester/server.

2. RPC will be used on top of TCP/IP as the communication vehicle.

3. We will define the concept of the Unit of Work.

4. We will do data conversion.

5. Security:

   - Use Kerberos for authentication
   - Let each database system handle its own security.

6. If the server has more than one database, we require that the databases be homogeneous.

Work to Be Done

Before the next meeting, we will find out:

- The effort of the OSI/RDA group related to this subject.
- The effort of the OSF group related to this subject.
- The effort of other vendors related to this subject.
- Advertising this Working Group.
- A first draft will be written.

Attendees

David Benton	benton@bio.nlm.nih.gov
Richard Bowles	bowles@stsci.edu
Robert Enger	enger@seka.scc.com
Elizabeth Feinler	
Charles Fineberg	fineberg@wums2.wustl.edu
Russ Hobby	rdhobby@ucdavis.edu
Mike Janson	mjanson@mot.com
Greg Minshall	minshall@wc.novell.com
Joel Replogle	replogle@ncsa.uiuc.edu
Paul Selkirk	paul@ftp.com
Daisy Shen	daisy@watson.ibm.com
Jim Sheridan	jsherida@ibm.com
Wing Fai Wong	wfwong@malta.sbi.com
Wengyik Yeong	yeongw@psi.com

3.1.7 Network Fax (netfax)

Charter

Chair(s):

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 facsimilies 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 maximize the possibility for interoperability among network fax projects.

Among the issues that need to be resolved are what actual protocol(s) 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 development of a gateway system that will allow existing Group 3 and Group 4 fax machines to operate in a network environment.

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 that some fax vendors, as well as the networking community and fax gateway developers, will be brought into the effort.

Goals and Milestones:

Done — Review and approve charter making any changes deemed necessary. Refine definition of scope of work to be accomplished and initial 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 i e 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.

CURRENT MEETING REPORT

Reported by Mark Needleman/U California

NETFAX Minutes

The Netfax Working Group meeting was held on March 14, 1991 at the IETF meeting in St. Louis. The major purpose of the meeting was to discuss a proposal submitted by ISI to define a file format for facsimilies using TIFF. The idea behind this is that the problem of using facsimilies on the internet needed to be broken down into two separate and discrete tasks. One of which would be to agree on a common file format, and separately to come to some agreement on a transport mechanism.

The thought was that once the file format was agreed on there would be enough common agreement in place so that some experimentation could begin. Once the experimentation had taken place there could be later agreement on what was the best transport mechanism (SMTP or FTP) or perhaps multiple mechanisms could be used depending on what made sense in a particular environment.

There was a lot of agreement with this basic concept. There was also a good bit of discussion on whether TIFF was the best mechanism to use as proposed in the ISI paper or whether something using ODA should be the mechanism. It was concluded that the group did not have enough expertise with ODA to make that evaluation and what was needed was a definition of what an ODA encoding for facsimilies would look like so it could be compared to the TIFF encoding as presented in the ISI paper. Peter Kirstein (P.Kirstein@cs.ucl.ac.uk) agreed to provide such a definition and post it to the list.

The ISI paper is available for anonymous ftp from stubbs.ucop.edu as:

/pub/netfax/isi-faxpaper

A discussion was held over the transport mechanism for actually moving facsimilies around the network and whether it should be SMTP or X.400. Dave Crocker discussed some of the recent happenings at the Internet Mail Extensions (SMTPEXT) Working Group and his feeling that some of the mail header extensions that would be needed should be happening in the near future. He also mentioned how these headers would be compatible with X.400 whenever possible. This gave the group the sense that the issue of X.400 versus SMTP was not all that important since anything done for SMTP would most likely be compatible with X.400 headers.

Carl Malamud (carl@malamud.com) agreed to look at what headers exit in both SMTP and X.400 that could define the type of information that might be needed to transmit facsimilies across the network. This will enable us to determine if new headers need to be asked for from the SMTP group. It was also mentioned that once agreement was gotten on the headers, a combination of transport mechanisms could be used including FTP. All that would be needed for FTP would be to add the proper headers at the beginning of the file to be sent to a fax server. This common definition of headers along with a common agreement on file format holds out the promise of maximum interoperability among fax servers that might be deployed.

Action Items

- Peter Kirstein will post to the list, a definition of what an ODA encoding would look like. This will enable the group to quickly decide on the relative merits of ODA versus TIFF and make a final decision by the next meeting.
- Carl Malamud will post to the list, his investigation of mail headers. This will enable the group to decide what extensions need to be defined, if any, or whether what already exists is usable.
- Mark Needleman agreed to separate the netfax mail archives into multiple files for ease of downloading.

Attendees

David Crocker	dcrocker@pa.dec.c
Steve Deering	deering@xerox.com
Ned Freed	net@ymir.claremont.edu
Robert Hagens	hagens@cs.wisc.edu
Russ Hobby	rdhobby@ucdavis.edu
Mike Janson	mjanson@mot.com
Kenneth Key	key@cs.utk.gdy
Peter Kirstein	kirstein@cs.ucl.ac.uk
Anders Klemets	klemets@cs.cmu.edu
Jim Knowles	jknowles@trident.arc.nasa.gov
Shelly Knueven	shelly@wugate.wustl.edu
Tracy LaQuey Parker	tracy@utexas.edu
E. Paul Love	loveep@sdsc.edu
Clifford Lynch	lynch@postgres.berkeley.edu
Carl Malamud	carl@malamud.com
Mark Needleman	mhn@stubbs.ucop.edu
Ursula Sinkewicz	sinkewic@decvax.dec.com
Wing Fai Wong	wfwong@malta.sbi.com
Wengyik Yeong	yeongw@psi.com

3.1.8 Network News Transport Protocol (nntp)

Charter

Chair(s):
Eliot Lear,

Mailing Lists:
General Discussion: `ietf-nntp@turbo.bio.net`
To Subscribe: `ietf-nntp-request@turbo.bio.net`

Description of Working Group:

This group will study and review the issues involved with netnews transport over the Internet. Originally released as an RFC in February of 1986, NNTP is one of the widest implementations of an elective status protocol. As of this writing, the protocol has just passed its fifth birthday, not having been updated once.

Over the years several enhancements have been suggested, and several have even been implemented widely. The intent of this working group will be to encode the more popular and plausable enhancements into an Internet standard. Included in the inital list of changes to be considered are the following:

o user level and site designated authentication methods;

o binary transfer capability;

o minimization of line turnaround; and

o stronger article selection capability.

It is expected that public domain software will be released concurrently with an RFC, demonstrating the protocol enhancements.

Goals and Milestones:

Mar 1991 Define scope of work.

Jun 1991 Submit internet draft for review and comment.

Jun 1991 Possibly meet at USENIX for further comment.

Jul 1991 Meet at IETF for further comment.

Aug 1991 Submit RFC to IESG.

3.1.9 Network Printing Protocol (npp)

Charter

Chair(s):
Glenn Trewitt, trewitt@pa.dec.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.
Done	Write draft LPR specification.
Done	Discuss and review the draft LPR specification. Discuss long-range printing issues in the Internet. Review status of Palladium print system at Project Athena.
Done	Submit final LPR specification including changes suggested at the May IETF. Discuss document on mailing list.
Done	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 Glenn Trewitt/DEC

NPP Minutes

Agenda

- LPR/LPD Protocol RFC
- Printer Access Protocol – modifications
- Son of LPR/LPD; Palladium
- Charter and Schedule
- Network Printing Protocol from UMD

This meeting was hampered by a lack of continuity. Only four out of the twenty people at the meeting had been to any previous meetings. Advance notice of the next meeting may help with this.

Printer Access Protocol

There were several discussions before the meeting with members of the Security Area Advisory Group (SAAG) about how to add security to PAP. John Linn, who sat in on the meeting, was most helpful. Surprisingly, we were able to come up with a small set of extensions that do security to everyone's satisfaction. A note will be sent out describing these.

There was no discussion about the other issues mentioned in the Agenda, because Ajay Kachrani and Glenn Trewitt were the only individuals who had specific knowledge of them. Glenn has not seen any comments about the proposed changes that he sent out, or about the use of (minimal) PDL commands for paper tray, font, etc., selection mentioned in the Agenda.

LPD Protocol RFC

There was a very useful discussion about the nit-picky things that the RFC isn't clear on, such as acknowledgements. A revised RFC will be sent out with these elaborations within two weeks. An attempt will be made to deal with the following issues that have been raised at previous meetings:

- "Pure protocol" vs. 4.2 implementation.
- Noting extensions that have been made.

It is possible that some of the useful (compatible) additions may make it into 4.4 bsd. This would be a big win.

Network Printing Protocol from UMD

Bruce Crabill from the University of Maryland presented a protocol used there for printing. It resembles SMTP, in the form of its client/server dialog. The functionality is a bit higher than LPR/LPD. The significant improvement over LPR/LPD is the fact that responses can be more detailed, and that information can be passed back to the client. (In LPR/LPD, the only way that information gets back to the client is at the end of communication, in which case a text string (usually an error message) is sent back.)

Son of LPR/LPD; Palladium

There are still a lot of ideas about what belongs in the client `->spooler` , spooler `->spooler` , and spooler `->printer` protocols. There seemed to be considerable agreement that the three had only minor differences between them. This would lead to the consideration that perhaps there should only be one protocol. Is PAP a candidate? What about the UMD work? Glenn wants to see some discussion about this on the list *before* the next meeting.

Network Printing Working Group Charter

There was no discussion of the Charter or schedule, although Glenn intends to have either PAP or the LPR RFC ready for a final round of comments by the next meeting, and the other polished up by the next one.

Attendees

Charles Bazaar	bazaar@emulex.com
Bruce Crabill	bruce@umdd.umd.edu
Bill Durham	durham@MDC.COM
Elizabeth Feinler	
Tom Grant	grant@xylogics.com
Keith Hacke	hacke@informatics.wustl.edu
Ajay Kachrani	kachrani@regent.enet.dec.com
Neil Katin	katin@eng.sun.com
Kenneth Key	key@cs.utk.gdy
Charles Kimber	
Anders Klemets	klemets@cs.cmu.edu
John Linn	linn@zendia.enet.dec.com
David Lippke	lippke@utdallas.edu
Joshua Littlefield	josh@cayman.com
Leo McLaughlin	ljm@ftp.com
Donald Merritt	don@brl.mil
Keith Moore	moore@cs.utk.edu
Michael Patton	map@lcs.mit.edu
Jan Michael Rynning	jmr@nada.kth.se
Sam Sjogren	sjogren@tgv.com
Glenn Trewitt	trewitt@pa.dec.com

3.1.10 TELNET (telnet)

Charter

Chair(s):
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 will examine RFC 854, "Telnet Protocol Specification", in light of the last 6 years of technical advancements, and will determine if it is still accurate with how the TELNET protocol is being used today. This group will also look at all the TELNET options, and decide which 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.

TELNET Minutes

Agenda

- Telnet Environment Option
- Telnet Authentication Option
- Telnet Encryption Option

The Telnet Working Group met the morning of Tuesday, March 12, 1991, and the afternoon of Wednesday, March 13, at the St. Louis IETF meeting.

Telnet Environment Option

The first item of discussion was the ENVIRON option. Vint Cerf was present to express some of the views of the IAB on this option, and their reluctance to endorse it.

The crux of the issue is the fact that the ENVIRON option allows for arbitrary environment variable information to be passed between systems and that the draft RFC has no well-defined variables in it, the lack of the latter causing even more concern about the former. Vint suggested that submitting the ENVIRON option with some well- defined variables, and without the unknown variables being allowed, unless there was some good justification, could expedite the IAB accepting the ENVIRON option.

A list was put together of what well-known variables should be in the initial draft: The list was USER (LOGNAME), JOB, ACCT, PRINTER, SYSTEMTYPE and XDISPLAY. Dave Borman will write up a description of the format of the values for these and send them to the mailing list for discussion.

Because there is a strong feeling that giving the user the ability to pass arbitrary environment variable information is very useful, discussion was held on how to continue. One item that needs to be taken care of is to identify how to differentiate between well-known variables and user-defined variables. One option was to encode the information in the variable name, for example, ala the X-foo naming used in mail. The other option was to add a new code, USERVAR, that would have the same semantics as VAR, but be explicitly for non-standard variable names. A vote was taken, with three options:

1. Encode information in name.
2. Add USERVAR.
3. Leave it out for now, and don't worry about it.

With seven votes recorded, three voted for adding USERVAR, one voted for encoding in the name, and three voted for leaving it out for now. Hence, any future discussion for dealing with user-defined variables will use the USERVAR code.

Dave Borman will look into Vint' suggestion that it might be good for someone to go to an IAB meeting and present the reasons for the user-defined variables.

Telnet Authentication Option

The Authentication option was next on the Agenda. The revised draft, with definitions for Kerberos Version 4, was discussed. It became apparent that the NAME subcommand in the Kerberos definitions was something that could be needed by many authentication schemes, so the NAME suboption was moved up to its own suboption:

IAC SB AUTHENTICATION NAME remote user IAC SE

Two new options for Kerberos were added, CHALLENGE and RESPONSE, to provide mutual authentication. After the server authenticates the client, the client sends the server a CHALLENGE, an eight octet value encrypted in the session key. The server decrypts it, adds one to it, re-encrypts it, and sends it back in a RESPONSE command. If the client can successfully decrypt it, and get the original challenge value plus 1, then the server has been authenticated to the client. As an additional step, both sides take the original encrypted challenge, and encrypt it again in the session key, and save that new value for a unique encryption key that can be used by the ENCRYPT option. Hence, the NEWKEY command isn't needed anymore, and was therefore removed. The ACCEPT command was modified to remove the optional "authenticated principal", as it provided no new, useful information. There was a bit of discussion about the difference between authentication and authorization. A user may be able to authenticate on the remote machine, but still not be authorized to log in as the user specified in the NAME suboption. Also, this knowledge might not be known to the telnet server. Hence, the Kerberos REJECT command may or may not contain an explanation, and the client might well get an ACCEPT command, only to then later see a failure message from some other part of the remote system that fails the authorization.

A decision was made that, with these changes, the authentication option is fairly stable. The changes will be incorporated into the document and distributed for review, and if there are no major objections it will be sent off to be published as an RFC. The Kerberos definitions will be removed and published as a separate document.

Telnet Encryption Option

One item of a rather lengthy discussion entailed the security aspects of the Encryption option. The net result was that it was decided that for now the document would state that the encryption option provides protection against a passive attacker (i.e, someone who is snooping in on the packets as they fly by), but not against an active attacker (i.e., someone who is snooping, and can also intercept/modify the packets as they fly by). The crux of the discussion was for when the encryption option is normally off, and is only being enabled in one direction when sensitive information is passing over the network, like passwords. Later versions of the option may contain information about how to provide adequate protection against an active attacker.

Key exchange was also discussed. In all cases, key exchange is currently outside of the scope of the Encryption option. It is assumed that there are one or more keys available that are known to each side of the connection. It was decided that the START and REQUEST-START would have an additional argument added, a keyid. A keyid is an arbitrary length number. It is encoded with the MSB first, and the LSB last. All the bytes between the START/REQUEST-START and the IAC SE are the keyid. A one-byte keyid with a value of zero was reserved to mean "the default key". This will usually refer to a key derived as a side effect of authentication. For all other keyids, an algorithm is needed to do the exchange of information to decide which key name to use. David Borman agreed to write something up on this.

> [Begin additional info, not part of the minutes of this meeting]
>
> What will be in the next draft is the addition of two new commands: ENC_KEYID and DEC_KEYID. The side that is going to encrypt sends ENC_KEYID with a keyid that it understands. The decrypting side responds with a DEC_KEYID command with the same value if it accepts it, a different value if it doesn't accept it but has a different keyid to try, or an empty value if there are no more values. If the encryptor receives a different value than what it sent, it processes it in the same way, sending over one of the three possible responses. This continues until both sides have sent and received (or received and sent) the same value.
>
> [End of additional info]

The initial description on Kerberos DES encryption that was in the latest draft document was modified quite a bit. It was decided that we needed a definition for both Cipher Feed Back (which is what was already documented, more or less...) and for Output Feed Back. The Initial Vector is sent by the encryptor, and is sent as a clear text string across the network. The view was that this was probably okay, but there was some concern that it might need to be encrypted. However, for now it will just be clear text. The encrypter sends across the IV, and the decryptor sends back either an IV_OK or IV_BAD message. If IV_OK is received, then negotiation of the keyid, happens, and then encryption can be enabled/disabled as needed.

The Telnet Working Group will meet at the July IETF Meeting in Atlanta.

Attendees

Richard Basch	probe@mit.edu
David Borman	dab@cray.com
Vinton Cerf	vcerf@NRI.Reston.VA.US
Tom Grant	grant@xylogics.com
Neil Haller	nmh@bellcore.com
Russ Hobby	rdhobby@ucdavis.edu
John Linn	linn@zendia.enet.dec.com
David Lippke	lippke@utdallas.edu
George Sanderson	sanderson@mdc.com
Jeffrey Schiller	jis@mit.edu
Sam Sjogren	sjogren@tgv.com
Dean Throop	throop@dg-rtp.dg.com
William Townsend	townsend@xylogics.com
Kathleen Wilde	wilde@decvax.dec.com

3.2 Internet Area

Director(s):

- Noel Chiappa: jnc@ptt.lcs.mit.edu

Area Summary

Reported by Greg Vaudreuil/CNRI

- IP over Appletalk
  Discussed a couple of issues. They are working on an Appletalk over IP tunneling specification and have a rough draft. Additional efforts include working on a variety of MIBs for Appletalk as well as working on IP over the Point to Point Protocol.

- Multi Media Bridging
  There have been proposals to change some facets of IP over 802 media to make bridging easier. A document proposing changes was written, however the proposal would incompatibly change the way IP is used over 802.5 networks. There was some general discussion on what to do with the concept of transparent bridging. The group felt that this idea, while not pure or attractive would not go away. Rather than let chaos reign, the group opted to list the problems with the transparent bridging, and attack the problems one at a time. Some of the issues include the hardware byte order, and fragmentation. One firm idea is... if the box is not a pure bridge, the box must be a full router, doing all the TTL, mtudiscovery, and other router functions.

- Router Requirements
  The Router Requirements Working Group has a new version of the document. The document is for the most part complete. The only missing pieces are network management and miscellaneous applications. Outstanding technical issues include protocol leaking between protocols, i.e., Rip to OSPF interactions, route pruning, and Type of Service issues. A standard mechanism for picking among routes is needed to avoid loops.

- Dynamic Host Configuration
  No information available at the time of the plenary report. See the Minutes included later in these Proceedings.

- Connection IP
  Finish ST2. ST2 is an interim protocol by which the group can experiment with resource allocation issues. This protocol will be run on the DARTNET testbed. There are implementations of the protocol for experimentation.

CURRENT MEETING REPORT

Reported by Andy Nicholson/Cray Research, Inc.

CBNR BOF Minutes

These are the Minutes from the "Conditioning of By-Request Network Resources" Birds of a Feather session which met at the St. Louis IETF. Due to the small size and informality of the meeting, no formal minutes were taken. This record is believed to be reasonably accurate and proper credit given to the originators of the ideas and concepts presented. My apologizes for any errors or omissions.

The meeting began with a short exposition from Andy Nicholson about the purpose of the meeting and some description of work done at Cray Research Inc., for the support of Circuit Switched T3 networks. While working with circuit-switched T3 networks, developers at Cray Research Inc., determined that there would be advantages to defining a standard way to control certain classes of network resources through the internet. In the case of a circuit-switched T3 line, the line should be switched on only when there are active transport connections which can fully utilize the service. Due to the high cost of the resource, underutilization would be particularly undesirable. The developers believe that this capability might have other applications in the internet and that an effort should be made to define a standard protocol. It was noted that this work involved a host on the internet sending internet messages to another host which communicated with a T3 switch, and could turn the switch on and off.

Dan Friedman offered the suggestion that a more refined architectural model could be used and that hosts would often be less concerned with accessing a particular network connection than with making a particular class of service available. He suggested that messages should be formatted to request an abstract service, rather than control a specific service provider directly.

Jeff Young and Andy Nicholson were both uncomfortable with this idea, as existing products do not exist to use this capability, and Cray Research was already working to provide a resource-specific allocation capability for interested customers. They felt that it was necessary to support direct access to specific resources.

Numerous discussions followed, during which Dan also noted that routing policy would be involved in decisions whether to allocate network resource. A four-layer architectural model emerged from these discussions:

- Policy Layer

  Handles policy questions like "Will I allocate a resource to satisfy this request from this requester?"

- Resource Layer

Makes decisions regarding which of many possible resources to allocate to satisfy a particular request.

- Action Layer

  Handles the mechanics of allocating a particular instance of a type of resource.

- Hardware Layer

  Actual network resources to be allocated and de-allocated.

In an actual system, each layer would be represented by some processing occurring on a host somewhere in the internet, except for the hardware layer which might not be capable of internet connectivity (i.e., a T3 circuit switch accessible only by a dialup line). When a resource is desired, a message would be sent to the "Policy Manager" (the entity residing at the policy layer), which would determine the disposition of the request.

In a real system, the Policy and Resource managers might be null, and simply pass requests on the layer below. This will allow the implementation of a system where a host makes direct requests for specific network resources (i.e., a specific T3 switch to connect two particular hosts).

It was also agreed that routing policy is being explored by another group, so we would not work on policy layer issues. Furthermore, we did not see an immediate need to work on resource layer issues. We agreed that since there is an immediate need to define an interface to the action layer, we would work on that. The interface between the action layer and the hardware layer is hardware-dependent, and will need to be implemented on a case-by-case basis. In the model, action layer direct messages would be sent to the policy layer, but neither the policy nor resource layers are yet defined and exist as null entities.

Some of the information that the action manager would require appeared obvious and was:

- Request type - what to do.
- Resource identifier - what to do it to.
- Status - probably a return value.
- Endpoints - parties using the allocated resource.

Jeff Young postulated that there might be some vendor-specific information associated with the allocation of a specific resource. Jeff felt that this information might best be stored with the entity requesting the service and that the vendor specific information be passed in the request message from the requester. Not all were thrilled with this idea and it was suggested that this information should be maintained by the action manager and that the resource identifier should be sufficient to find any vendor-specific information that might be required to allocate the resource.

It was also suggested that there might be accounting information in the request messages, but it was noted that this might not always be necessary. It was also suggested that only the policy and/or resource managers would be interested in this information and that it should not be propagated to the action manager.

The vendor-specific data and accounting information issues got a lot of discussion, and it was suggested that we could define a message option format, much like tcp or ip options. In addition, we could pre-define at least two option types, vendor-specific data and accounting information. This idea was not universally popular either. If we meet at the next IETF (as the Chair hopes), these issues will require further discussion.

In the closing minutes of the meeting (it should be noted that we met on two consecutive nights), we came up with some additional details. We would put the address of the intended manager into the request messages. If the manager receiving a message is not the intended recipient, then that manager will forward the message (as in the case of a policy manager receiving action manager messages).

We also considered the possibility of a hierarchical message format, wherein the core message is an action manager message, and resource and policy information are added to the core message format, depending on the granularity of the requester's request. This was not decided at this meeting.

Dan Freidman and Andy Nicholson agreed to do some work on an RFC to document the protocol the group is working on.

If the interested parties are able, we will meet at the next IETF meeting.

Attendees

David Borman	dab@cray.com
Daniel Friedman	danfriedman@bbn.com
Joseph Golio	golio@cray.com
Andy Nicholson	droid@cray.com
Jeff Young	jsy@cray.com

3.2.1 Connection IP (cip)

Charter

Chair(s):
Claudio Topolcic, topolcic@nri.reston.va.us

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:

Done	Produce a new specification of ST.
Done	Define common hardware and software platform.
Done	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/CNRI

CIP Minutes

Agenda

- Status reports
  - COIP-K
  - ST-II
  - VT and PVP
  - SRI activities
- Discussion
  - Analysis of COIP approach vs other CL approaches

Meeting Report

Guru, Claudio and Steve gave overviews of the status of the implementations that they are responsible for.

Barbara gave an overview of the activities at SRI.

- Benchmarks on DARTnet.
- SFQ (based on source & destination IP addresses only) - implemented but not debugged.
- SFQ + resource reservation - to work with ST, for example.
- Writing an annotated bibliography on congestion control.
- tg currently uses tcp or udp sockets; we need to add ST sockets and test. Benchmark results: BW, loss, delay; fairness, path utilization.

Discussion of CO vs. CL Approaches

The purpose of this discussion was to understand the real differences between the approach taken by this group versus other, ostensibly connectionless, approaches that have been proposed, and where there are differences, to identify analysis, measurements, or experiments that would give us a better understanding of which approach is superior in which situation.

Steve led a discussion of our understanding of an alternate CL approach. The following is a diagram of the modules that would have to be implemented in a router in order to support such an approach.

```
 ------------------------------------------------------------
 |      ----------------       -------------       --------  |
 |     |     Packet    |      | Resource    |             | |
------> | Identification |-----> | Enforcement |----> Queue |----->
 |     |_______________|      |_____________|      --------| |
 |             |          -----        ^                     |
 |             |         |     |       |                     |
 |             |         |_____|       |                     |
 |             --------> |_____|--------                     |
 |                       |     | Forwarding                  |
 |                       |_____| Table                       |
 |                          ^                                |
 ------------------------- | ---------------------------------
                            |
                    ---------------
                   |               |
                   |   Resource    |
                   |   Manager     |
                   |               |
                   |_______________|
```

We discussed what were believed to be differences in the approaches.

1. Classes vs. individual flows.

   A proposed CL approach may have "classes" that can carry traffic belonging to different flows. However, Guru's MCHIP protocol has PICons and Lixia's Flow Protocol (FP) has Flow 0, either of which can carry packets from any flow so are equivalent concepts. When you use a PICon, you have to include more addressing info than just the logical channel number, perhaps the full addresses. This raised the question of whether the short headers that ST and MCHIP use are worthwhile, and how often they would be used?

   We may have a different view of the future. Will individual flows be small or large with respect to available bandwidth. If they are large, then identifying individual flows will be more important. If they are small, then perhaps it is better to aggregate a number of flows together. The answer may be different if we look at the short term or the long term.

2. Reservation request and the start of the data flow.

   There may be a difference as to the chronological binding of reservation to the time flow begins. We make the reservation at the time the flow begins. An alterate approach might allow a reservation ahead of time. There are some further issues, specifically, if the intent is to not do any work at the time the flow begins, then the system must be prepared to redo work as the topology changes.

3. Failure recovery.

   When a link goes down, connectionless protocols can reroute more easily if multiple paths exist. But in the CO scheme, we could use Flow 0 or PICon (or encapsulate ST in IP) along the alternate path without guarantees during the recovery. How fast will IP rerouting be compared to CO connection repair? One RTT?

4. Location of resource manager.

   The alternate approach allows the resource manager to be in a separate box from the router. A resource manager separate from the router allows a hot standby for redundancy, possibly fewer resource managers than forwarders, allowing the use of dumb, and therefore cheap, forwarders, and may simplify the transition from the current IP to an "integrated services" IP since the changes to the routers might be less so it would be easier to get the vendor to accept the change.

   However, it needs a reliable protocol between the resource manager and the forwarder, which must be standardized to allow mixing vendors and introduces a number trade-offs, e.g., problems because the manager doesn't directly see connectivity changes. Further, we don't expect any difference in setup time required with separate resource manager vs. one combined in the router.

5. Transition path to the new system.

   A CL approach is presumed to allow an easier transition. However, how significant is it whether the first 20 bytes look the same as an IP header? In either case, new software must be installed in all routers that need to implement resource management. Host software may not need to change if resource management used only IP options since the existing BSD software allows IP options to be specified by the application.

6. Resource management.

   This is an issue regardless of the approach taken. Furthermore, in general, the same mechanisms can be used in both approaches.

7. Flywheel resource allocation.

   This is a scheme by which a router predicts the resource requirements of flows within a implicitly by monitoring past usage and assuming that the requirements will change slowly, that is, it has "momentum". If a new flow is detected which would overuse a class's resources, that new flow could be blocked. This approach requires keep-alives, may require further feedback to the applications, and does not interact well with pre-scheduling of resources.

8. Routing.

   A CO oriented approach doesn't need smart routing because the routes are verified anyway, allows for alternate path routing based on load whereas a datagram approach does not, because it is unstable. Further, we couldn't see how IP multicast would support dynamic flows efficiently.

9. Explicit vs implicit setup.

   A CO scheme, which naturally incorporates explicit setup, allows coordinated call blocking, which would allow for some set of related flows to succeed, rather than a random set. However, in an implicit setup scheme, the cost (delay) is the same if the setup fails, but much lower if it succeeds, which is presumed to be most of the time. On the other hand, doesn't just push the buck up a level (making the application decide if connection didn't work, vs. having explicit setup at a lower layer)?

Experiments

We identified a number of tests and experiments that could be conducted to try to tell which approach may be better under what circumstances.

- Questions
  - Does blocking work?
  - How much interference comes from outages?
  - Do you honor scheduled calls?
  - Utilization?
- Types of experiments:
  - Measure lost bandwidth due to flywheel approach as utilization approaches saturation.
  - If CO implies enforcement per flow, and CL allows enforcement per class, which works better.
  - Failure recovery.
    * What is the impact of an outage on flows over paths that haven't failed (as failed flows are rerouted)?
    * How long does it take to reconstruct and what mechanisms are required in each case?
    * Measure time required to detect failure with various schemes.
- What is the setup time?
- How well are pre-scheduled flows honored?
- Flip-side of (1): How much loss due to momentum of the flywheel (time the allocation is held after the flow stops) and what is the impact of reducing the timeout?
- Which approach is better for correlated flows?

Attendees

Joe Blackmon	blackmon@ncsa.uiuc.edu
Andreas Bovopoulos	andreas@patti.wustl.edu
Helen Bowns	hbowns@bbn.com
Stephen Casner	casner@isi.edu
Barbara Denny	denny@sri.com
Zubin Dittia	zubin@dworkin.wustl.edu
Allison Mankin	mankin@gateway.mitre.org
Jay Melvin	infopath@well.sf.ca.us
Gary Mussar	mussar@bnr.ca
Andy Nicholson	droid@cray.com
Philippe Park	ppark@bbn.com
Gurudatta Parulkar	guru@flora.wustl.edu
Rehmi Post	rehmi@ftp.com
K.K. Ramakrishnan	rama@kalvi.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
Brad Solomon	bsolomon@hobbes.msfc.nasa.gov
Martha Steenstrup	msteenst@bbn.com
Claudio Topolcic	topolcic@NRI.reston.va.us
Wing Fai Wong	wfwong@malta.sbi.com

3.2.2 Dynamic Host Configuration (dhc)

Charter

Chair(s):

Ralph Droms, droms@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 be automated (i.e., those that must be managed by network administrators).

Goals and Milestones:

Done	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.
Done	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.
TBD	Write a bootp extensions document

CURRENT MEETING REPORT

Reported by Ralph Droms/Bucknell

DHC Minutes

The discussion at this meeting was driven by three primary Agenda items:

- BOOTP forwarding agent document – Walt Wimer
- Dynamic IP assignment protocol – Jesse Walker
- DHCP Internet Draft – Ralph Droms

There was also discussion of future work.

Walt Wimer prepared a detailed description of the BOOTP forwarding agent (which is only hinted at in the BOOTP RFCs) for use in the Router Requirements RFC and in the DHCP RFC. The Working Group decided the appropriate course of action would be to publish Walt's document as a separate RFC updating the original BOOTP RFCs [RFC-951, RFC-1084], with a reference from the Router Requirements RFC to this new RFC. The Working Group also discussed some changes and filled in some details in the new RFC. Walt is working on incorporating the Working Group's suggestions and some other clarifications to the original BOOTP RFCs to prepare his document for publication as an Internet Draft. A draft version of the revised BOOTP forwarding agent document is available for anonymous FTP from host `sol.bucknell.edu` in file `dhcwg/bootp-forwarding`.

Next, the Working Group discussed Jesse Walker's description of the dynamic IP address allocation and configuration parameter transmission algorithm. The Working Group was in general agreement with the description of the client-server protocol. There was a spirited discussion for and against the use of multiple DHCP exchanges for the transmission of configuration parameters, e.g., in the case where there are more parameters than could be transmitted in a single DHCP packet. This discussion interacted with an earlier discussion about negotiation for transmission of parameters: a client may need to request certain, specific parameters while a server may need to send parameters that were not requested but should have non-default values in the client. The "Tastes Great" contingent felt that restricting the client to a single DHCP request was too restricting, while the "Less Filling" contingent argued for simplicity and pointed to extension mechanisms (reusing fields in the BOOTP protocol specification, using TFTP to download larger configuration files) that could be used in those cases where the parameters could not all fit in a single DHCP packet. The Working Group concluded that it was likely that whoever wrote the protocol specification would get to decide the issue.

Two documents will be put up for consideration as Internet Drafts before the next meeting of the Working Group: Walt's BOOTP relaying agent document, and a description of the client-server component of DHCP, based on Jesse's contribution. At the next meeting, we will take up the server-server DHCP protocol. We must also begin discussion of an SNMP interface to DHCP; anyone anxious to write a MIB definition for DHCP?

There is a mailing list for this Working Group at host-conf@sol.bucknell.edu (administrivia to host-conf-request). An archive of the mailing list and other documents of interest are available for anonymous FTP from sol.bucknell.edu under directory dhcwg.

Attendees

Steve Alexander	stevea@i88.isc.com
Karl Auerbach	karl@eng.sun.com
Richard Basch	probe@mit.edu
Lida Carrier	lida@apple.com
Ralph Droms	droms@bucknell.edu
Karen Frisa	karen.frisa@andrew.cmu.edu
Robert Gilligan	gilligan@sun.com
Tom Grant	grant@xylogics.com
Ajay Kachrani	kachrani@regent.enet.dec.com
Tom Kessler	kessler@sun.com
Jim Kinder	jdk@fibercom.com
Darren Kinley	kinley@crim.ca
Joshua Littlefield	josh@cayman.com
Gary Malkin	gmalkin@ftp.com
Bill Manning	bmanning@houston.sc.ti.com
Leo McLaughlin	ljm@ftp.com
Linda Melvin	infopath@well.sf.ca.us
Greg Minshall	minshall@wc.novell.com
Robert Morgan	morgan@jessica.stanford.edu
Brad Parker	brad@cayman.com
Michael Patton	map@lcs.mit.edu
Jonathan Saperia	saperia@decwrl.enet.dec.com
Tim Seaver	tas@mcnc.org
Paul Selkirk	paul@ftp.com
Glenn Trewitt	trewitt@pa.dec.com
Mike Turico	mturico@mot.com
John Veizades	veizades@apple.com
Kathleen Wilde	wilde@decvax.dec.com
Walter Wimer	walter.wimer@andrew.cmu.edu
John Wobus	jmwobus@suvm.acs.syr.edu
Wing Fai Wong	wfwong@malta.sbi.com

3.2.3 IP over Appletalk (appleip)

Charter

Chair(s):
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:

Done	Describe, in an RFC, the current set of protocols used to connect Macintoshes to IP internets.
Done	Define a MIB for the management of DDP/IP gateways.

INTERIM MEETING REPORT

Reported by John Veizades/Apple

APPLEIP Minutes

The Apple-IP Working Group met on January 9th in conjuction with the San Francisco MacWorld Exposition. This is a summary of the decisions that were made at that meeting.

The MacIP and AppleTalk MIB are both approaching standards submission.

There was some discusion about the standardization of the AA protocol but no conclusions were made.

The standardization of AppleTalk PPP extensions will be brought up for discussion at the St. Louis IETF meeting. Brad Parker and Frank Slaughter will work on a proposal on this protocol.

The AppleTalk Tunneling document was discussed for the rest of the meeting. The following open topics were left.

- The authority over the registration of UIDs Apple was proposed.
- What to do with hop count.
- How to maintain the tunneling gateways routing tables.

The group came to the conclusion that implementation was possible and developers were encouraged to do so.

The next meeting will be at the St. Louis IETF meeting with subsequent meetings at the following possible times:

- Apple Developers Conference May 91
- Atlanta IETF August 91 or
- MacWorld Boston August 91
- InterOp Oct 91
- Los Alamos IETF Nov 91

CURRENT MEETING REPORT

Reported by John Veizades/Apple

APPLEIP Minutes

AppleTalk over IP Tunneling

Alan Oppenheimer made several comments on modifications he has made to the document to cover some operational experience they have acquired in their development process.

Issues still to be resolved include Zone name explosions, security and cross-router coordination and routing.

MIBs and SNMP

The MIB draft document is in the Internet-Draft directory and several implementations are in progress or completed (Cayman, Shiva, Apple, etc.).

The Working Group should begin thinking of defining a specification of running SNMP over DDP (AppleTalk Network Protocol).

The AppleTalk MIB II will be discussed at the next meeting. Steve Waldbusser is working with Apple to define a specification for end node management of AppleTalk hosts.

PPP and Atalk

Frank Slaughter (Shiva) is working toward a specification of this protocol.

AA Protocol

Phil Bunde from Shiva is working toward a specification of this.

MacIP

The MacIP document can be found in the Internet-Draft directory. Comments are strongly solicited by the author. Implementation of both client and server must be worked on.

The next meeting of this Working Group will be during the week of the Apple Developers Conference the week of May 13. The meeting will be a full or half day meeting with the possibility of an AURP interoperability workshop.

Attendees

Peter Caswell	uunet!pacvax!pfc
Karen Frisa	karen.frisa@andrew.cmu.edu
Fred Gray	fred@homer.msfc.nasa.gov
Douglas Kerr	dougk@wc.novell.com
Mary Louise Laier	laierl@applelink.apple.com
Nik Langrind	nik@shiva.com
Steven Lendt	network@zeus.unomaha.edu
Joshua Littlefield	josh@cayman.com
Milo Medin	medin@nsipo.nasa.gov
Greg Minshall	minshall@wc.novell.com
Robert Morgan	morgan@jessica.stanford.edu
Alan Oppenheimer	oppenheimer1@applelink.apple.com
Brad Parker	brad@cayman.com
Robert Peglar	robp@anubis.network.com
Evan Solley	solley@applelink.apple.com
Brad Solomon	bsolomon@hobbes.msfc.nasa.gov
Natee Tanchanpongs	natu@wugate.wustl.edu
John Veizades	veizades@apple.com
Jonathan Wenocur	jhw@shiva.com

3.2.4 IP over FDDI (fddi)

Charter

Chair(s):
Dave Katz, dmk@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:

Done 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.

3.2.5 Multi-Media Bridging (mmb)

Charter

Chair(s):
Jeffrey Fitzgerald, `jjf@fibercom.com`

Mailing Lists:
General Discussion: `mmbwg@fibercom.com`
To Subscribe: `mmbwg-request@fibercom.com`

Description of Working Group:

The Multi-Media Bridge Working Group has the task of addressing the function of multi-media bridges within TCP/IP networks. This is viewed as necessary at this time because of the proliferation of these devices.

The first goal of the group is to document the multi-media bridge technology and point out the issues raised by having these devices in a TCP/IP internet. If there are problems which can be addressed the group will work towards resolving them and documenting the solutions.

Goals and Milestones:

Done	Finalize Charter of Group
Aug 1991	Document mulit-media bridging technology and its affect on TCP/IP Internets.
Aug 1991	Document issues to be addressed by Working Group.

CURRENT MEETING REPORT

Reported by Jeff Fitzgerald/Fibercom, Inc.

MMB Minutes

Report not submitted.

3.2.6 Point-to-Point Protocol Extentions (pppext)

Charter

Chair(s):

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 designed 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. In addition it will define new PPP options for the existing protocol definitions, such as stronger authentication and encryption methods.

Goals and Milestones:

none specified

3.2.7 Router Discovery (rdisc)

Charter

Chair(s):

Steve Deering, deering@xerox.com

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 than 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.

3.2.8 Router Requirements (rreq)

Charter

Chair(s):

James 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.

The purposes of this project include:

- Defining what an IP router does in sufficient detail that routers from different vendors are truly interoperable.
- Providing guidance to vendors, implementors, and purchasers of IP routers.

The requirements developed will be split into two volumes. The first will cover link layer protocols and address resolution. The second will cover everything else. We intend that the link layer protocol document will apply not only to routers but also to hosts and other IP entities.

The Working Group will also instigate, review, or (if appropriate) produce additional RFC's on related topics.

Goals and Milestones:

Done	First Internet Draft version of the upper layer volume.
Oct 1990	First Internet Draft version of the link layer volume.
Dec 1990	Second Internet Draft version of upper layer volume.
Dec 1990	Second Internet Draft version of link layer volume.
Feb 1991	Third Internet Draft version of upper layer volume.
Feb 1991	Third Internet Draft version of link layer volume.

CURRENT MEETING REPORT

Reported by Philip Almquist/Consultant

RREQ Minutes

Shortly before the St. Louis meeting, the second Internet Draft version of the Router Requirements specification was released. Except for a few open issues (described below), the technical content of the document is fairly well set. Extensive editorial work remains to be done.

On the first day of the meeting, the Chair conducted a brief "Introduction to Router Requirements" session for first-time attendees and anyone else who was interested. After that, the Working Group dove into four half days of meetings.

Three of the four sessions were devoted to fine-tuning the draft. Particular attention was paid to chapters 3 (Link Layer), 4 (Internet Layer Protocols), 9 (Miscellaneous Application Layer Protocols), and 10 (Operations and Maintenance). We also discussed what still needed to be done to complete the draft. Items identified included:

- Chapter 8 (Network Management) still needs to be written.
- Much of chapter 9 (Miscellaneous Application Protocols) still needs to be written.
- Coverage of security-related topics needs to be extended.
- There should be additional discussion sections providing implementation hints and explaining the rationale behind some of the requirements.
- Several smaller sections need to be written or revised.
- As mentioned above, extensive editorial work is still required.

Volunteers were solicited to do the necessary work. The issue of variable length subnet masks was also raised. Since the IETF Working Group on this topic is still not underway, we tentatively decided that Router Requirements would have to partially skirt this issue, though we will say more than the current draft does.

The remaining session was devoted to discussion of three important and interrelated issues:

1. Route choice -- how a router chooses which route to use to use for a packet when the router has several routes (learned from different routing domains) to the packet's destination.

2. Route leaking -- how a router which is in multiple routing domains decides whether a route learned in one routing domain ought to be advertised into other routing domains.

3. Route filtering -- how a router decides whether to disbelieve certain routes from certain sources.

Each of these issues has two components:

1. Constraints -- what must (or must not) be done to avoid undesirable phenomena such as routing loops and black holes?

2. Controls -- what sorts of configuration options does a network manager need to be able to do to make routing work in moderately complex parts of the Internet?

The Working Group was not able to reach any consensus on these issues, but will continue to try to address them in the time before the IETF meeting in Atlanta in July. However, several Working Group members also participated in a productive Border Gateway Protocol (BGP) Working Group session which addressed issues specific to route leaking between BGP and OSPF.

Frank Solensky deserves considerable commendation for diligently noting all of the changes to the draft which were agreed to during the course of the meeting.

Attendees

Philip Almquist	almquist@jessica.stanford.edu
Douglas Bagnall	bagnall_d@apollo.hp.com
Atul Bansal	bansal@netrix.nac.dec.com
Ballard Bare	bare@hprnd.rose.hp.com
William Barns	barns@gateway.mitre.org
Bob Beach	bob@ultra.com
Arthur Berggreen	art@acc.com
Helen Bowns	hbowns@bbn.com
David Bridgham	dab@asylum.sf.ca.us
Ronald Broersma	ron@nosc.mil
Christopher Bucci	bucci@pluto.dss.com
Jeffrey Burgan	jeff@nsipo.nasa.gov
Ross Callon	callon@bigfut.enet.dec.com
Graham Cobb	cobb@marvin.enet.dec.com
Rob Coltun	rcoltun@trantor.umd.edu
Don Coolidge	coolidge@speaker.wpd.sgi.com
Nabil Damouny	
Kurt Dobbins	dobbins@ctron.com
Dino Farinacci	dino@3com.com
Jeffrey Fitzgerald	jjf@fibercom.com
Richard Fox	rfox@synoptics.com
Vince Fuller	vaf@Standford.EDU
Fred Gray	fred@homer.msfc.nasa.gov
Martin Gray	mg@spider.co.uk
Jeremy Greene	greene@coral.com
Martin Gross	gross@polaris.dca.mil

Patrick Heisinger	heisinger@mdcguy.mdc.com
Jeffrey Honig	jch@devvax.tn.cornell.edu
Kathleen Huber	khuber@bbn.com
Harold Jones	hjones@nac.enet.dec.com
Ajay Kachrani	kachrani@regent.enet.dec.com
Frank Kastenholz	kasten@asherah.clearpoint.com
Manu Kaycee	kaycee@trlian.enet.dec.com
Stev Knowles	stev@ftp.com
Alex Koifman	akoifman@bbn.com
Anthony Lauck	lauck@tl.enet.dec.com
Eliot Lear	lear@net.bio.net
Tony Li	tli@cisco.com
Mike Little	little@ctt.bellcore.com
Joshua Littlefield	josh@cayman.com
Then Liu	
Gary Malkin	gmalkin@ftp.com
Mike Marcinkevicz	mdm@calstate.edu
Glenn McGregor	ghm@merit.edu
Milo Medin	medin@nsipo.nasa.gov
Carol Melowitz	melowitz@mdcgwy.mdc.com
Jay Melvin	infopath@well.sf.ca.us
Linda Melvin	infopath@well.sf.ca.us
April Merrill	
David Miller	dtm@ulana.mitre.org
Greg Minshall	minshall@wc.novell.com
John Moy	jmoy@proteon.com
Gary Mussar	mussar@bnr.ca
David O'Leary	oleary@sura.net
Brad Parker	brad@cayman.com
Stephanie Price	price@cmc.com
Michael Reilly	reilly@pa.dec.com
Kary Robertson	kr@concord.com.kr
George Sanderson	sanderson@mdc.com
Mark Schaefer	schaefer@davidsys.com
Steven Sherry	shsherry@eng.xyplex.com
Stephen Shew	sdshew@bnr.ca
Frank Solensky	solensky@clearpoint.com
Evan Solley	solley@applelink.apple.com
Brad Solomon	bsolomon@hobbes.msfc.nasa.gov
Michael St. Johns	stjohns@umd5.umd.edu
Martha Steenstrup	msteenst@bbn.com
Roxanne Streeter	streeter@nsipo.nasa.gov
Glenn Trewitt	trewitt@pa.dec.com
Walter Wimer	walter.wimer@andrew.cmu.edu
Cathy Wittbrodt	cjw@nersc.gov

John Wobus	jmwobus@suvm.acs.syr.edu
Wing Fai Wong	wfwong@malta.sbi.com
Chin Yuan	cxyuan@pacbell.com

3.2.9 Service Location Protocol (svrloc)

Charter

Chair(s):
John Veizades, `veizades@apple.com`

Mailing Lists:
General Discussion: `srv-location@apple.com`
To Subscribe: `srv-location-request@apple.com`

Description of Working Group:

The Service Location Working Group is chartered to investigate protocols to find and bind to service entities in a distributed internetworked environment. Issues that must be addressed are how such a protocol would interoperate with existing directory based services location protocols. Protocols that would be designed by this group would be viewed as an adjunct to directory service protocols. These protocols would be able to provide a bridge between directory services and current schemes for service location.

The nature of the services location problem is investigative in principle. There is no mandate that a protocol should be drafted as part of this process. It is the mandate of this group to understand the operation of services location and then determine the correct action in their view whether it be to use current protocols to suggest a services location architecture or to design a new protocol to compliment current architectures.

Goals and Milestones:

Done	Open discussion and determine if a Working Group should be formed.
Mar 1991	Continue discussion trying to refine the problem statement and possible resolutions.
Jul 1991	Do we take the RFC track or do we write a report on our conclusion and leave it at that?

CURRENT MEETING REPORT

Reported by John Veizades/Apple

SVRLOC Minutes

The Service Location Protocol group came to a consensus that the work that was being done was of the Working Group direction and that the group should convene as a Working Group and not as a BOF.

Work is being done on the same type of protocols in the IRTF by Michael Schwartz. John Veizades said he would pursue understanding how that work could be leveraged by the group.

An overview of the result of the Boulder BOF was presented.

The group brainstormed on ideas that might solve some of the problems that were represented and came to the conclusion that there is a gap between the services provided by directory services and the type of protocols that were discussed at this meeting.

As a statement of architectural direction is needed by the group, Leo McGlaughlin, Steve Waldbusser and John Veizades will meet before the next meeting to try to firm the Charter up and come to some terms with the architectural direction.

Attendees

Karl Auerbach	karl@eng.sun.com
Richard Basch	probe@mit.edu
David Bridgham	dab@asylum.sf.ca.us
Randy Butler	rbutler@ncsa.uiuc.edu
Lida Carrier	lida@apple.com
Bill Durham	durham@MDC.COM
Fred Engel	engel@concord.com
Karen Frisa	karen.frisa@andrew.cmu.edu
Tom Grant	grant@xylogics.com
Keith Hacke	hacke@informatics.wustl.edu
Russ Hobby	rdhobby@ucdavis.edu
Harold Jones	hjones@nac.enet.dec.com
Ajay Kachrani	kachrani@regent.enet.dec.com
Douglas Kerr	dougk@wc.novell.com
Darren Kinley	kinley@crim.ca
Stev Knowles	stev@ftp.com
Mary Louise Laier	laierl@applelink.apple.com
Mike Little	little@ctt.bellcore.com
Joshua Littlefield	josh@cayman.com
Leo McLaughlin	ljm@ftp.com

Greg Minshall	minshall@wc.novell.com
Agnes Moran	
Brad Parker	brad@cayman.com
William Townsend	townsend@xylogics.com
Glenn Trewitt	trewitt@pa.dec.com
Mike Turico	mturico@mot.com
John Veizades	veizades@apple.com
Steven Waldbusser	steven.waldbusser@andrew.cmu.edu
Jonathan Wenocur	jhw@shiva.com
Walter Wimer	walter.wimer@andrew.cmu.edu

3.2.10 Special Host Requirements (shr)

Charter

Chair(s):

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 truly 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:

Done Mailing list discussion of Charter and collection of concerns.

Done	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.

3.3 Network Management Area

Director(s):

- James Davin: jrd@ptt.lcs.mit.edu

Area Summary reported by James Davin/MIT

A number of Working Groups met at the St. Louis IETF meeting. Brief summaries of their activities are presented below. More detailed accounts are presented in the Minutes for each Working Group.

Also, at the St. Louis meeting, the SNMP Network Management Directorate met and considered five items of business.

1. The Ethernet-like Interfaces MIB, as amended by the SNMP Working Group at St. Louis, was discussed and reviewed positively.

2. The IP over AppleTalk MIB produced by the IP over Appletalk Working Group was discussed and reviewed positively.

3. The state of the SMDS Interface Protocol MIB introduced to the SNMP Working Group at St. Louis was discussed.

4. The current draft of the OSPF MIB produced by the OSPF Working Group was given a preliminary review in anticipation of its imminent completion. A list of comments was provided to the author.

5. A document describing enhancements to the SNMP administrative model to provide better support for security and proxy configurations was discussed at length and reviewed positively after amendments designed to minimize change to existing infrastructure. The updated version of this document was subsequently distributed for wider review to the SNMP Security Working Group.

Internet Accounting (acct)

The Internet Accounting Working Group met at St. Louis and focused on definition of MIB objects for the collection of accounting information. They also spent some time on minor revisions of the accounting architecture document, which will be distributed soon as an Internet Draft. The Accounting Working Group is also coordinating with the Remote LAN Monitoring and Operational Statistics Working Group to assure that no redundant MIB instrumentation is defined.

Bridge MIB (bridge)

The Bridge MIB Working Group met and, with specific amendments, recommended their consensus document for consideration as a Proposed Standard. One strictly informational question will be resolved via electronic mail. The approved Bridge MIB text will be distributed by electronic mail for working members to verify the agreed amendments. The revised text will be available before the next meeting.

During the Bridge MIB Working Group meeting, and also in the IETF Plenary session, presentations were made to clarify the IETF policy on the translation of network management definitions developed by other standards bodies into the SNMP idiom. A letter that addresses the particular case of the Bridge MIB effort has been incorporated into the Bridge MIB Working Group Minutes. Its three enumerated points capture the general policy for "importing" MIBs adopted by the IESG.

Character MIB (charmib)

The Character MIB Working Group met and reviewed the current working documents. These were recommended by the Working Group for consideration as Proposed Standards with specific amendments. Revised text reflecting these amendments will be available soon.

FDDI MIB (fddimib)

The FDDI MIB Working Group met and accomplished all goals set at its previous meeting. Work on defining instrumentation was completed, as was work on mapping actions in the ANSI specifications into appropriate SNMP MIB objects. The need for 64-bit integer MIB objects was obviated, and a document defining two SNMP traps to model FDDI events was introduced at this meeting.

Management Services Interface (msi)

The Management Services Interface Working Group met briefly at the St. Louis meeting and adjourned owing to low attendance and the absence of key individuals. The Working Group affirmed its decision that the interface is only relevant to management stations. Those assembled reviewed the list of outstanding issues and found two particularly problematic: (a) the translation between the OSI and SNMP information models and (b) the tension between the usefulness of an opaquely defined interface and the need to offer guidance to implementors.

DECNet Phase IV MIB (decnetiv)

The DECNet Phase IV MIB Working Group made many changes to the current document including significant reduction in the number of objects. The definition of events was relegated to a distinct document, generic portions of the X.25 instrumentation were excised, and the definition of conformance groups was revised. The Working Group reached consensus on the text as amended. Revised text will be available within four weeks.

Remote Lan Monitoring (rlanmib)

The Remote LAN Monitoring Working Group met and reviewed most of the current draft in detail. A revised draft will be posted to the mailing list soon. An interim Working Group meeting will be held in 4–6 weeks to continue discussion.

SNMP

The SNMP Working Group met and considered two documents. The Working Group recommended the Ethernet-like Interfaces MIB, with very minor amendments, for consideration by the IESG as a Proposed Standard. The approved text will be posted to the Working Group mailing list soon after the meeting. A document defining a SMDS Interface Protocol MIB was introduced to the Working Group in a presentation made by Kaj Tesink of Bellcore. The Working Group will consider this document further in subsequent meetings.

3.3.1 Bridge MIB (bridge)

Charter

Chair(s):
Fred Baker, fbaker@emerald.acc.com

Mailing Lists:
General Discussion: bridge-mib@nsl.dec.com
To Subscribe: bridge-mib-request@nsl.dec.com

Description of Working Group:

The Bridge MIB Working Group is a subgroup of the SNMP Working Group, and is responsible for providing a set of SNMP/CMOT managed objects which IEEE 802.1 Bridge Vendors can and will implement to allow a workstation to manage a single bridged domain. This set of objects should be largely compliant with (and even drawn from) IEEE 802.1(b), although there is no requirement that any specific object be present or absent.

Goals and Milestones:

Done	Publish initial proposal
Done	Submit an Internet Draft
Feb 1991	Submit draft for RFC publication

CURRENT MEETING REPORT

Reported by Fred Baker/ACC

BRIDGE Minutes

The Bridge MIB Working Group convened for two sessions on Tuesday, March 12, 1991. The Bridge MIB under review had been posted to the bridge-mib discussion group on February 16, 1991, as Working Document Bridge MIB, Draft 6, by Decker, Langille, Rijsinghani, and McCloghrie.

Chair, Fred Baker opened the meeting with a review of the proposed Agenda, which had been posted to the discussion group earlier. All present agreed to the Agenda, which follows:

- Static Table (dot1dStatic)
  Four objects in the entry
- Window Table (dot1dWindow)
  Two objects in the entry
- Base Group/Port Table (dot1dBase)
  Three objects in the group
  Three objects in the port entry
- STP Group/Port Table (dot1dStp)
  Fourteen objects in the group
  Eleven objects in the entry
- SR Group/Port Table (dot1dSr)
  Sixteen objects
- Transparent Group/FDB and Port Tables (dot1dTp)
  Two objects in the group
  Three objects in the FDB Entry
  Five objects in the Port Entry

Anil Rijsinghani presented the dot1dStatic table. He gave a review of how Forwarding Database (FDB) entries come to be and how entries in the static Table are made. The SNMP concept of entries with a status of "invalid" was also discussed.

Jim Kinder initiated a lengthy discussion of the relationship of table entries with a dot1dStaticReceivePort of value zero to other entries. Various hypothetical scenarios were discussed and the resulting decision was that an entry with a dot1dStaticReceivePort can coexist along with other entries for the same address.

The discussion of the dot1dStatic table was suspended when the Working Group was addressed by Phill Gross, IETF Chair. He talked about the interaction between the IETF and the IEEE 802.1d Committee. Last year a letter was received from the IEEE expressing concern about a possible duplication of efforts by the IETF Bridge MIB Working Group and the IEEE 802.1d Committee.

Phill reviewed the IETF philosophy for using the work of a standards body in conjunction with its own work. The IETF will use the reference work as a starting point, while being free to subset it, and within the confines of sound engineering principles, to augment it.

A draft of a response letter to the IEEE was presented (see below) and the group approved of sending it along with a copy of the Bridge MIB.

Jeff Case pointed out that we need to be sensitive to the fact that a reference document that is used for a starting point may change as work is done within the IETF and that an incompatibility may result between the final reference document and the final IETF work.

After the break, talk resumed on the the dot1dStatic table. The agreement was that an entry in the table with dot1dStaticReceivePort=0 is the default value to use if a specific dot1dStaticReceivePort is not specified.

The hierarchy of the Forwarding Database is this, then.

- Static information for a specific receive port (dot1dStaticReceivePort¿0).
- Static information for all ports (dot1dStaticReceivePort=0).
- Learned information.

The dot1dStatic table was approved with wording to accomplish the above changes.

Keith McCloghrie presented the dot1dTpFdbWindowTable, starting with an overview of the design considerations

The Problem: To provide an efficient means of retrieving the whole or a significant portion of a transparent bridge's Forwarding Database.

Alternatives:

- Get-Next Sweep - 1 Powerful Get Next per Conceptual Row
  1 Conceptual Row per Round Trip

- Bulk Algorithm (RFC 1187)
  either - 1+ Powerful Get Next per Conceptual Row
  say, 3 Conceptual Rows per Round Trip

- or say, 1 Powerful Get Next per 2 Conceptual Rows +
  1 Get Request per 10 Conceptual Rows

- say, 4 Conceptual Rows per Round Trip

- Window Table: 1 Powerful Get Next per 40 Conceptual Rows
  40 Conceptual Rows per Round Trip

Advantages:

- Less ASN.1 encoding/decoding (size and performance)
- Can access starting in the middle of a table (e.g., all DECNET addresses)

Disadvantages:

- Have to look into data to get address for next Powerful Get Next
- DUMB MIB sweepers will retrieve redundant information. (but in the same number of requests)

Why 40?

- A round number
- PDU size < 576
- Benefit of > 40 not considered worth the effort

Keith then compared the dot1dTpFdbTable and the dot1dTpFdbWindowTable, noting that they contain the same number of entries.

```
    Window Table                              FDB Table

         ---------------                           ---------------
         |      N      |                           |      N      |
        _|_____________|                          _|_____________|
        |  (N-1),N    ||                          |     N-1     ||
        |             |                           |             |
       .             .                           .             .
      .             .                           .             .
     .             .                           .             .
    ---------------                           ---------------
    |    2-41     |                           |      2      |
   _|_____________|                          _|_____________|
   |    1-40     ||                          |      1      ||
   |             |                           |             |
   |_____________|                           |_____________|
```

A discussion of the dot1dTpFdbWindowTable followed Keith's presentation.

Bob Stewart argued for including 42 entries from the FDB in each dot1dTpFdbWindowTable. He presented a sound engineering underpinning for his argument but the group decided to leave the number at 40.

A corollary discussion took place about the viability of having a variable length window. Jeff Case pointed out that the SNMP Protocol Specification says in part:

"An implementation of this protocol need not accept messages whose length exceeds 484 octets."

He recommended that the Bridge MIB should not allow arbitrarily large PDUs. The Working Group agreed to leave the number at 40.

A question was raised about the dot1dTpFdbWindowTable being in the spirit of SNMP vis a vis, not supporting aggregate objects. Jeff Case spoke once again and indicated that although the dot1dTpFdbWindowTable did not particularly excite him, he had no philosophical objection to it.

Various optimization ideas were presented for Powerful Get Next walks and although no consensus was reached, four options were discussed for the disposition of this table.

1. Delete it.
2. Leave it in the MIB as is (Status Optional).
3. Port it to another document to be developed in the Experimental tree.

4. Leave it in the MIB, but change the status to Mandatory.

No consensus was reached and the dot1dTpFDBWindow group discussion was tabled.

After the break, Chair Fred Baker led a review of the document section by section.

Keith McCloghrie clarified that the wording "protocols that are bridged" is used to differentiate between those PDUs that are bridged versus those that are not.

Bill Anderson spoke from a user's perspective. He presented a need for the Bridge MIB FDB to cover all addresses, bridged and otherwise. Various members of the group pointed out that the Remote LAN Monitoring group was addressing this issue, and in fact had specified this functionality.

Two IEEE 802.1d managed objects were left out of the "not included" group on page 8. These are SpanningTreeProtocolPort objects DiscardLackOfBuffers and DiscardErrorDetails. These will be added.

The discussion moved to the dot1dBase group.

Bob Stewart noted that bit ordering for the "Bridge ID" was not specified, and it was necessary here and other places in the document.

The discussion moved to the dot1dStp group.

The incompatibility between IEEE 802.1d specification of time in 1/256ths of a second and the Bridge MIB of 1/100ths of a second was brought up. A challenge was issued by Fred Baker to name a chip that gave 1/256ths granularity for its clock, and the issued died. A side issue of the syntax of dot1dStoMaxAge was brought up. After a discussion of the correct use of TimeTicks vs INTEGER, no change was recommended.

A change was made to dot1dStpPriority description to uniquely identify two octets within the Bridge ID.

Maurice Turcotte pointed out that dot1dStpPortMulticastAddress should not be on a per port basis and that this address can be determined from the variables dot1dStpProtocolSpecification and ifType. This variable was included at the request of Eric Decker and since he was not present, the group decided to delete this variable, and allow Eric to comment.

In the afternoon session, the broken(6) dot1dStpPortState was discussed at great length. No agreement was reached and the issue was tabled.

Steve Sherry requested that new TCN counters be added. The consensus of the group was that these counters would present information available elsewhere and were most useful for debugging code rather than networks. No variables were added for TCN counters.

A discussion of BridgeID vs. (Priority - Address) with respect to dot1dStpPortDesignatedPort. The broad issue was whether to represent BridgeID variables as one variable or separated into BridgePriority and BridgeAddress. The decision was to leave the variables as they are in the document.

The range of dot1dStpPortPathCost should have been (1-65535)

The dot1dSr group passed without comment.

The dot1dTp group passed without comment.

After a brief recess the broken(6) dot1dStpPortState was revisited.

The two major points raised in favor of keeping this state were:

1. We need to know when the Spanning Tree Protocol cannot bridge through a port because it is dysfunctional and it would be nice to know that from one variable and,

2. It is possible for the Spanning Tree Protocol to have the port in forwarding state and the port be non-operational.

The two major points against this state were:

1. There is no broken(6) state in the Spanning Tree Protocol and,

2. This information is already available from the combination of ifAdminStatus, ifOperStatus, and dot1dStpPortState.

After more intense discussion, the group reached consensus and removed the broken(6) value from the dot1dStpPortState.

Next the dot1dFdbWindow group was reopened for discussion. After a brief discussion, the consensus was reached that we separate the dot1dFdbWindow group into a separate document and develop it further in the Experimental branch of the MIB.

Next the traps were reviewed and agreement was reached after some discussion to let the traps stand. A slight modification was made to the newRoot trap description, with a view to ensuring (to the extent possible) that the Network Management Station would be able to receive the trap.

The Bridge MIB Working Group agreed to forward the Bridge MIB Draft 6, with the above modifications, to the IETF for acceptance as a Proposed Standard.

LETTER TO IEEE 802.1

William P. Lidinsky
Chairman, IEEE 802.1
Institute of Electrical and Electronics Engineers

Dear Mr. Lidinsky:

Enclosed with this letter, please find the current working draft of the SNMP Bridge MIB, produced by the IETF Bridge MIB Working Group.

The Bridge MIB Working Group was organized under the IETF's Network Management Directorate in May 1990, and has studied the semantics of 802.1(d) with a goal of representing it in an SNMP SMI-compliant MIB.

The IETF wishes to cooperate with, and coordinate its MIB development efforts with, other ongoing MIB development activities in other standards organizations. In cases where the IETF wishes to develop an SNMP MIB for technology already being considered by another standards group, we have established the following policy:

1) The IETF will always utilize the current effort of another group as the starting point for its own MIB development activities. Therefore, a major portion of the IETF effort may simply be translating the other MIB into the SNMP SMI idiom.

2) Because the requirements of other organizations may not be precisely the same as those of the IETF, we may choose initially to include only a subset of the other MIB. In such a case, we would reserve the opportunity to consider adding the remaining objects to the SNMP MIB in the future.

3) In some cases, we may wish to propose additional objects based on operational experience. It is not expected that this would be a very common occurrence, and in such cases we would make every effort to communicate the IETF proposed objects back to the appropriate group for their consideration.

A comparison of 802.1(d) and the current IETF draft should show that, in fact, there are few significant differences.

I hope your group will have the opportunity to review the IETF SNMP Bridge MIB. We would appreciate hearing any comments or suggestions you may have.

We look forward to working together with you in the future.

Thank you,

IAB Chair IETF Chair IETF NM Area Director IETF Bridge MIB Working Group Chair

Attendees

William Anderson	wda@mitre.org
Fred Baker	fbaker@emerald.acc.com
Steve Bostock	steveb@novell.com
Howard Brown	brown@ctron.com
Theodore Brunner	tob@thumper.bellcore.com
Jeffrey Case	case@cs.utk.edu
John Casey	
Chris Chiotasso	chris@roswell.spartacus.com
Bill Durham	durham@MDC.COM
Nadya El-Afandi	nadya@network.com
Richard Fox	rfox@synoptics.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Martin Gray	mg@spider.co.uk
Jeremy Greene	greene@coral.com
B.V. Jagadeesh	bvj@esd.3com.com
Frank Kastenholz	kasten@asherah.clearpoint.com
Manu Kaycee	kaycee@trlian.enet.dec.com
Kenneth Key	key@cs.utk.gdy
Jim Kinder	jdk@fibercom.com
Cheryl Krupczak	clefor@secola.columbia.ncr.com
Then Liu	
Keith McCloghrie	kzm@hls.com
John Morrison	jfm@lanl.gov
Robert Peglar	robp@anubis.network.com
David Perkins	dave_perkins@3com.com
Anil Rijsinghani	anil@levers.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
Dror Shindelman	pbrenner@sparta.spartacus.com
Bob Stewart	rlstewart@eng.xyplex.com
Maurice Turcotte	dnmrt@rm1.uucp

3.3.2 Character MIB (charmib)

Charter

Chair(s):
Bob Stewart, `rlstewart@eng.xyplex.com`

Mailing Lists:
General Discussion: `char-mib@decwrl.dec.com`
To Subscribe: `char-mib-request@decwrl.dec.com`

Description of Working Group:

The Character MIB Working Group is chartered to define an experimental MIB for character stream ports that attach to such devices as terminals and printers.

The Working Group must first decide what it covers and what terminology to use. The initial thought was to handle terminals for terminal servers. This directly generalizes to terminals on any host. From there, it is a relatively close step to include printers, both serial and parallel. It also seems reasonable to go beyond ASCII terminals and include others, such as 3270. All of this results in the suggestion that the topic is character stream ports.

An important model to define is how character ports relate to network interfaces. Some (a minority) terminal ports can easily become network interfaces by running SLIP, and may slip between those states.

Given the basic models, the group must select a set of common objects of interest and use to a network manager responsible for character devices

Since the goal is an experimental MIB, it may be possible to agree on a document in 3 to 9 months. Most of the group's business can be conducted over the Internet through email.

Goals and Milestones:

Done	Mailing list discussion of Charter and collection of concerns.
Done	Discuss and final approval of charter; discussion on models and terminology. Make writing assignments.
Done	First draft document, discussion, additional drafts, special meeting?
Dec 1991	Review latest draft and if OK, give to IESG for publication as RFC.

CURRENT MEETING REPORT

Reported by Bob Stewart/Xyplex

CHARMIB Minutes

The Character MIB Working Group held its fourth meeting at the IETF meeting in St. Louis. Attendance is falling off a bit as we near completion, but we had representatives of several terminal server companies and other interested parties. Overall, the meeting showed good consensus and resulted in the completion of the business at hand. As a result, following edits based on the meeting and a final implementation report, we are ready to submit our documents to become proposed standards.

The following meeting Agenda was presented informally on the Character MIB Working Group mailing list before the meeting.

Agenda

- Discuss the drafts as distributed via the mailing list.
  - Character MIB
  - RS-232-like MIB
  - Parallel-printer-like MIB
- Discuss implementations.
- Recommend drafts for advancement to proposed standard.

The group pointed out minor editorial errors in all three documents.

We discussed the necessity of common values for types of hardware flow control and decided to let simple practicality win over architectural purity, adding values for CTS/RTS and DTR/DSR to the Character MIB flow control types.

After considerable discussion about including flow-control characters in the Character MIB's port character counts, we decided to leave them included but qualify the counts to be "detected" characters.

We discussed and approved the overall changes adding a synchronous group to the RS-232 MIB. We decided to differentiate between frames short-terminated due to aborts and those interrupted by modem signal changes.

We discussed the implementation and operational cost of signal change counts in the RS-232 and Parallel MIBs and decided to leave them in pending an implementation report.

We agreed that all objects with "Err" in the name should have "Errs" instead.

The only implementation in progress is being done by Xyplex for inclusion in a product this year. It is near completion and will provide the first implementation test.

Pending satisfactory edits and the completed implementation, the group agreed to recommend all three documents for advancement as proposed standards.

Attendees

Charles Bazaar	bazaar@emulex.com
Christopher Bucci	bucci@pluto.dss.com
Shawn Gallagher	gallagher@quiver.enet.dec.com
Tom Grant	grant@xylogics.com
Frank Kastenholz	kasten@asherah.clearpoint.com
Kenneth Key	key@cs.utk.gdy
Jim Kinder	jdk@fibercom.com
David Perkins	dave_perkins@3com.com
Bob Stewart	rlstewart@eng.xyplex.com
Dean Throop	throop@dg-rtp.dg.com

3.3.3 DECnet Phase IV MIB (decnetiv)

Charter

Chair(s):

Jonathan Saperia, saperia@decwrl.enet.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 assignments will be made.

Done — Mailing list discussion of Charter and collection of concerns.

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 1991 — Make document an Internet Draft. Continue revisions based on comments received at meeting and over e-mail. Begin 'real' implementations.

Mar 1991 — 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

DECNETIV Minutes

The meeting in St. Louis was devoted to editing the DECNet Phase IV MIB Document. Important issues resolved were:

- Adjustment of several of the conformance groups.
- Removal of the X.25 and Event groups from the next draft.
- The addition of a level 1 routing table the routing group.
- The addition of Pointers to the IfIndex to several of the groups.
- Agreement on a number of typographic corrections and editorial changes which will appear in the next draft.

Agreement was reached to incorporate approximately 40 changes in the draft and then to send the revised document to the Working Group for review. After this final review it is hoped to have the document posted as an Internet Draft.

In addition to the items above, there were a few action items:

- Dean Throop will report to the group on his X.25 findings and status.
- Steve Hunter will determine if LineMaxBlockSize is the MTU.
- Jon Saperia will investigate if a counter timer object is needed for all relevant groups.

Attendees

Chris Chiotasso	chris@roswell.spartacus.com
Fred Engel	engel@concord.com
Mike Erlinger	mike@mti.com
Charles Fineberg	fineberg@wums2.wustl.edu
Steven Hunter	hunter@es.net
Jim Kinder	jdk@fibercom.com
Jay Melvin	infopath@well.sf.ca.us
David Perkins	dave_perkins@3com.com
Jonathan Saperia	saperia@decwrl.enet.dec.com
Dror Shindelman	pbrenner@sparta.spartacus.com

3.3.4 FDDI MIB (fddimib)

Charter

Chair(s):
Jeffrey Case, `case@cs.utk.edu`

Mailing Lists:
General Discussion: `fddi-mib@CS.UTK.EDU`
To Subscribe: `fddi-mib-request@CS.UTK.EDU`

Description of Working Group:

The primary goal of the FDDI MIB Working Group is to define a MIB for FDDI devices with objects which are based on those defined in the ANSI FDDI specifications and are compliant with the Internet standard SMI, MIB, and SNMP.

Goals and Milestones:

Sep 1990	"Final" initial draft of required get/set variables.
Oct 1990	Initial implementations of required get/set variables.
Feb 1991	Revised "final" draft of required get/set variables.
Mar 1991	Adoption of draft of required get/set variables.
Mar 1991	Initial draft of traps (events) and actions.

CURRENT MEETING REPORT

Reported by Jeff Case/UTenn

FDDIMIB Minutes

The following is an expanded version of the Minutes which appeared in the 19th Proceedings of the Internet Engineering Task Force, (Boulder, CO).

The FDDI MIB Working Group met on Wednesday, December 5, 1990. The meeting was held in conjunction with the IETF plenary at NCAR, Boulder, Colorado.

The items covered were as follows:

- Welcome and Introductions

  Once again, there were several attendees who were new to the Working Group. The goals of the group were reviewed for their benefit.

- Status Reports

  - Document. The text of the current document draft was distributed and discussed. The draft can be retrieved via anonymous FTP from /pub/fddimib/fddimib.txt at cs.utk.edu. Minutes and Agendas may be found there as well.
  - ANSI. Our ANSI counterparts in X3T9.5 SMT held their meeting at the same time as the plenary. Unfortunately, their meeting was in California and ours was in Colorado. This made attendance at both meetings difficult, but not impossible. The IETF FDDI meeting was schedule so as to be on the "lightest" day of the ANSI schedule at the request of those who desired to participate in both meetings. James Reeves and Bert Williams reported on the ANSI meetings held earlier in the week. They are working through the ballots to try to resolve the comments in hopes there will be a consensus on the next ballot.
  - Implementations. Several independent and interoperable implementations of the specification exist. Several others are underway. The restrictions associated with unannounced products made this discussion necessarily vague.

- Changes in the current issue of the document were discussed.

  Several changes were made based on problems identified during the implementation process.

  - The subranges on several integers was corrected from 1..65536 to 1..65535
  - The snmpFddi prefix was added to all variables which would have had a name space collision on the OBJECT DESCRIPTOR with the ANSI name(s)
  - The "-" was removed from all variable names. Implementation experience

showed that some user interfaces, when faced with a variable name such as snmpFddiSMTT-Notify would attempt to subtract one unknown variable (Notify) from another unknown variable (snmpFddiSMTT) resulting in chaos.

  - The first draft of the text for definitions of MIB variables to implement the SNMP equivalent of SMT Actions was added.

- Issues
  - Paths were discussed. PATHs are severely broken. The Working Group decided to delete all variables in the PATH group unless and until ANSI fixes them.
  - Groupings. The groupings of the variables were discussed and found to be satisfactory.
  - Optional variables. It was decided to delete all optional variables or to make them mandatory. The chair was tasked with preparing appropriate text. The Working Group considered what to do with the optional variables which are not in the document. Two major positions were taken. One position was that some vendors are going to implement the optional variables and the existence of a standard document will ease interoperability because they would all implement them the same way, that is, the argument was for the preparation and publication of a document such as the one prepared by Fox and Williams after a similar discussion at the Vancouver meeting. The contrary position was that the very existence of a document will create market forces which will lead to requirements that vendors implement all of the optional variables, even if their usefulness is questionable and their status as optional often meant that consensus could not be met in the ANSI community. After considerable discussion, it was decided by the group that the optional variables would not be worked on further and that they would not be published. It was further decided that this could be revisited when and if the signals coming from the ANSI community become more clear. That is, the optional variables might be considered when the document goes from Proposed to Draft. [N.B.: The line above was incorrect in the minutes originally distributed ... it read Draft to Proposed. I have never been able to get this right. JDC]
  - Traps. Traps were discussed. Since traps are somewhat contentious and the Concise Trap document does not enjoy the same status as the Concise MIB document, it was decided that it would be wise to decouple the MIB definitions and the TRAP definitions, lest we stall progress on the MIB definitions. The TRAP document is to be discussed at the next meeting.

- Action

  The Working Group voted to recommend the issuance of the MIB document as an RFC, pending final review of the text. This is to occur in early 1991.

The primary technical output of the meeting resulting from the review of the current draft was a decision to restructure the variable groups so as to allow a single network application entity (agent) to support more than one SMT.

Future work will entail:

1. review and comment on the mandatory get/set variables defined thus far,
2. Gaining implementation experience with the above,
3. Engineering ANSI events and actions into traps and MIB variables in accord with Internet standards, and
4. Addressing optional groups.

Bert Williams (Synernetics) and Rich Fox (SynOptics) volunteered to work on the text for the optional variables and to forward them to the Chair for inclusion in the draft at an appropriate time.

Current Draft

The text of the current draft may be obtained via ftp from anonymous(guest)/pub/fddimib/fddi-mib.txt at cs.utk.edu.

Next Meeting

The next meeting of the FDDI MIB Working Group is tentatively scheduled to be held in in conjunction with Interop '90. The meeting will most likely be held on Thursday evening. The primary topic of discussion will be to review implementation experiences and interoperability issues uncovered in the preparation for and performance of the Interop event. Plans for the meeting will be announced via the mailing list as they are finalized.

Attendees

Alan Apt	76307.3176@compuserve.com
Jack Brown	jbrown@huachuca-emh8.army.mil
Jeffrey Case	case@cs.utk.edu
Cho Chang	chang_c@apollo.hp.com
Chris Chiotasso	chris@roswell.spartacus.com
Paul Ciarfella	ciarfella@levers.enet.dec.com
Burt Cyr	burt@uncng.com
James (Chuck) Davin	jrd@ptt.lcs.mit.edu
Nadya El-Afandi	nadya@network.com
Richard Fox	sytek!rfox@sun.com
Debbie Futcher	dfutche@relay.nswc.navy.mil
Scott Hiles	whiles@relay.nswc.navy.mil
Satish Joshi	sjoshi@synoptics.com
Frank Kastenholz	kasten@interlan.com
Shimshon Kaufman	
Jim Kinder	jdk@fibercom.com

Christopher Kolb	kolb@psi.com
Cheryl Krupczak	clefor@secola.columbia.ncr.com
Peter Lin	lin@eng.vitalink.com
Keith McCloghrie	kzm@hls.com
Donna McMaster	mcmaster@davidsys.com
David Perkins	dave_perkins@3com.com
James Reeves	jreeves@synoptics.com
Anil Rijsinghani	anil@levers.enet.dec.com
Marshall Rose	mrose@psi.com
Jonathan Saperia	saperia@tcpjon.enet.dec.com
Jeffrey Schiller	jis@mit.edu
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Dean Throop	throop@dg-rtp.dg.com
Bert Williams	bert.synernetics@mailgate.synnet.com
Jeff Young	jsy@cray.com

CURRENT MEETING REPORT

Reported by Jeff Case/UTenn

FDDIMIB Minutes

The FDDI MIB Working Group last met on Monday afternoon, March 11, 1991, in St. Louis, Missouri. The meeting was held in conjunction with the Twentieth meeting of the Internet Engineering Task Force.

The Minutes of the previous meeting (Boulder) were adopted as distributed with one correction: the discussion of the decision to forego consideration of the optional variables should read "... the optional variables might be considered when the document goes from Proposed to Draft." The Minutes as originally distributed read "...Draft to Proposed."

The meeting began with a status report. As usual, there were several attendees who were new to the Working Group. The goals of the group were reviewed for their benefit. Version 0.9 of "FDDI Management Information Base" is nearing completion. Version 0.1 of "SNMP Trap Definitions for FDDI Management" is new. There was a brief discussion of the status of various early implementations of previous versions of the draft.

The members of the Working Group who also participate in counterpart efforts within ANSI reported on that groups activities and progress. The group is continuing to make progress toward resolving the changes resulting from comments received with the last round of letter ballots and converting the ANSI document to conform with the ISO standard "SMI" format known in that community as GDMO.

Most of the meeting was spent in detailed discussion of Version 0.9 of the MIB document. The net result was that the editor is to make a number of changes, the most important of which include the following:

1. Fix broken text in the description of the differences between version 0.9 and the previous version (or delete the revision history entirely);
2. Identify the version number of the ANSI SMT document which relates to the MIB document through an appropriate reference;
3. Place the appropriate text found in sister documents (like the dotx series) into Section 5, Overview;
4. Investigate renaming fddi OBJECT IDENTIFIER ::= experimental 8 to snmpFddi OBJECT IDENTIFIER ::= experimental 8 in the interest of consistency; (Editor's note: this can be done more easily when the subtree gets "promoted" to the transmission subtree)
5. Delete all references to the textual convention FddiTimerTwosComplement, replacing all such references to INTEGER;

6. Delete all references to the textual convention FddiTimeStamp, since, due to changes enumerated below, there are no longer any variables which utilize it;
7. Change the naming of several variables in order to obtain consistency – the first example of that is the name of smtNumber becomes snmpFddiSMTNumber – all variables will now begin with the prefix snmpFddi;
8. Add new explanatory text to snmpFddiSMTConnectionPolicy;
9. Delete snmpFddiSMTMsgTimeStamp and snmpFddiSMTTransitionTimeStamp and renumber snmpFddiSMTStationAction as a result;
10. Clone some clarifying, explanatory text from snmpFddiMACUpstreamNbr into snmpFddiMACOldUpstreamNbr;
11. Add a new enumeration of unknown(5) to snmpFddiMACDownstreamPORTType;
12. Change snmpFddiMACTMax, snmpFddiMACTMin, and snmpFddiMACTvxValue from read-write to read-only;
13. Delete the second, incorrect MAC from snmpFddiMACFrameMACCondition in two places in the document;
14. Create new variables corresponding to ANSI fddiMACFrameErrorThreshold, fddiMACFrameErrorRatio, and fddiMACUnaDaFlag;
15. Correct the semantics associated with snmpFddiATTACHMENTIMaxExpiration;
16. Define object identifiers for chipset types; and
17. Add several new contributors to the Acknowledgements section.

While this list is long, these changes are relatively minor when viewed in the scope of the history of the document. The length is mainly a result of the detailed description of the changes. The most significant change is the deletion of 64 bit counters and the removal of all the associated ugliness.

The editor was asked to make a decision and to pen appropriate text if necessary in response to a rather lengthy discussion regarding "detected" and "best effort" counters.

The editor was directed to make the above changes and submit the document for review and comment by the broader community through submission and publication as an Internet Draft. Following a brief period of review and comment to insure that the editor implemented the Working Group's wishes correctly, (e.g., a few weeks), it is the group's desire that unless there is widespread divisive discussion, the document be recommended for publication as an RFC as a Proposed Standard.

The Trap document was not discussed due to the lack of available time.

The last Agenda item was the status of the variables listed as optional in the ANSI document and therefore not included in the SNMP FDDI MIB document. The Working Group decided in Boulder that the optional variables would not be worked on further and that they would not be published at this time. It was decided that this could be revisited when and if the signals coming from the ANSI community become more clear, perhaps as the document goes from Proposed to Draft. This issue was placed on the Agenda because there were some individuals who disagreed with the result. There was no interest expressed by the attendees in reversing the decision at this time.

The current text of the MIB and TRAP documents may be obtained via anonymous ftp as:

- ~anonymous(guest)/pub/fddimib/fddi-mib.txt at cs.utk.edu and
- ~anonymous(guest)/pub/fddimib/trap.txt at cs.utk.edu.

Attendees

Howard Brown	brown@ctron.com
Jeffrey Case	case@cs.utk.edu
Cho Chang	chang_c@apollo.hp.com
Chris Chiotasso	chris@roswell.spartacus.com
Paul Ciarfella	ciarfella@levers.enet.dec.com
Don Coolidge	coolidge@speaker.wpd.sgi.com
Nabil Damouny	
Nadya El-Afandi	nadya@network.com
Joseph Golio	golio@cray.com
Jeremy Greene	greene@coral.com
Jim Kinder	jdk@fibercom.com
Cheryl Krupczak	clefor@secola.columbia.ncr.com
Then Liu	
John LoVerso	loverso@westford.ccur.com
Ron Mackey	
Keith McCloghrie	kzm@hls.com
David Perkins	dave_perkins@3com.com
James Reeves	jreeves@synoptics.com
Greg Satz	satz@cisco.com
Dror Shindelman	pbrenner@sparta.spartacus.com
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Dean Throop	throop@dg-rtp.dg.com
Mike Turico	mturico@mot.com
Bert Williams	bert@mailgate.synnet.com
Mark Wood	
Jeff Young	jsy@cray.com

3.3.5 Internet Accounting (acct)

Charter

Chair(s):
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:

Done 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

ACCT Minutes

The Internet Accounting Working Group met to:

- Review the results of the February meeting in Boston.
  - SNMP security and performance issues.

    SNMP seems a reasonable approach for transporting data, given a diskless meter, although FTP or other bulk file transfer mechanisms should also be allowed for meters which store accounting data on local disks. Other transport mechanisms may be discussed later.

  - Background Document.

    The background document can be released for general comment as an Internet Draft after the addition of PICTURES and explanations which illustrate how the accounting mechanism addresses a variety of scenarios. It is anticipated that the Background Document will be expanded again later.

  - Architecture Document.

    The existing architecture document can be released for general comment after revision and the addition of control parameters. Before it is released to the Internet Drafts area it will be posted to the Working Group mailing list for review.

  - Meter Services and MIB.

    The February discussion of control parameters and reporting formats was summarized for continuation.

- Discuss control parameters and reporting formats.
  - A modified reporting format resulted for further discussion.
  - A set of control functions was developed for further discussion.
  - The notion of being able to account differently on different interfaces and make finer distinctions resulted in the further development of a rule tree similar to

those discussed in February.

  - The ability to set the granularity of reporting in great detail through the use of a rule table was developed for further discussion. The current scheme seems too complex to be readily implemented, but serves as a starting point for further work. One solution to bounding the problem is defining a short list of standard (static) rule tables, without allowing the more general case.

A rough outline of the reporting format, control functions, rule tree, and rule table culled from the meeting notes and slides follows these Minutes.

- Additional notes about lengthy discussions.
  - It was noted that the ADDRESS_ID described in the reporting format might be expanded to transport level and beyond (e.g., application level), allowing for a more generalized accounting for any protocol stack, but that is beyond the charter of this Working Group.
  - It was also noted that attributes might be included in the ADDRESS_ID rather than as a separate field of the FLOW_ID.
  - Each packet shall be counted in ONE and ONLY ONE accounting record to avoid duplicate counts. Accounting records may be combined by the collection host for additional aggregate traffic information. This is a tentative response to the question Can a single packet be counted in multiple buckets of a single meter?
  - Meters in routers have special properties, since they are privy to the routing decision. Meters may be modelled as (a) one meter per interface (as a passive listener to the interface, not privy to the routing decision) or (b) one meter per router, aware of the both input and output interfaces for the packet. Passive listening devices must have a network address and possibly a separate connection to the network in order to be managed. Should routers be modelled as having a single meter to avoid complicating management?

Action Items

Background	Add pictures to Internet Background and revise. If changes are not too substantial, post directly to Internet Drafts.
Architecture	Revise Architecture Document to reflect control requirements and reporting changes. Post to Working Group mailing list for (time-limited) review before posting to Internet Drafts.

Meter Services	Make a stab at reducing the granularity control (rule table) problem to a manageable level. Further specify control parameters with the goal of creating a MIB.
Co-ordinate	Coordinate with Remote LAN MIB and Operational Statistics Working Goups since they may be tackling similar problems of granularity control.

REPORTING FORMAT (notes from discussion, not a precision representation):

```
Accounting Record ::=[ Meter ID and Unique Address provided by SNMP ]
    Start Time TIMESTAMP, [ optional ? ]
    End Time TIMESTAMP, [ should be current time ? ]
    Rule_Table ID? [ something might be needed here ...]
    SEQUENCE OF FLOW_RECORD. [ number of records, followed by records ]

FLOW_RECORD::=
    Flow FLOW_ID,
    Values VALUES.

FLOW_ID::=
    [0] Source ADDRESS_ID [ must have source or destination ]
    [1] Destination ADDRESS_ID or both ]
    [2] Subscriber_ID ADDRESS_ID [ optional ]
    [ Attributes not defined yet ]

VALUES::= [ rolling counters ]
    Fragments_Sent COUNT,
    Fragments_Rcvd COUNT, [ packets in the reverse direction are counted
                            here to avoid maintaining two accounting
                            records for a communicating pair - shouldn't
                            this be optional for source or destination
                            only flow ids? ]
    Bytes_Sent COUNT,
    Bytes_Rcvd COUNT, [ byte count of reverse flow ]
    First_Time TIMESTAMP, [ time first packet in flow seen if different
                            from meter start time ]
    Last_Time TIMESTAMP. [ time last packet in flow seen if different
                            from meter stop time ]

ADDRESS_ID::= [ some fields may be null, i.e., don't care ]
    [1] INTERFACE_INDEX INTEGER, [ as defined by SNMP ]
    [2] LINK LEVEL ADDRESS NETWORK_ADDRESS,
    [3] INTERNET ADDRESS NETWORK_ADDRESS,
```

[0] STRING OF OCTETS. [anything else used as unique ID].

NETWORK_ADDRESS ::=
Choice of {
[1] IP ADDRESS. (TCP/IP)
[2] NSAP ADDRESS. (OSI) variable length.
[n] X.25 Address (CCITT)
[m] MAC (LLC)
[x] STRING OF OCTETS. (any other arbitrary address) }

COUNT::=
Extensible_Integer SEQUENCE OF OCTETS.

TIMESTAMP ::= [defined by SNMP already, either absolute time or ticks/seconds/since meter boot time]

CONTROL PARAMETERS (notes under discussion):

Meter to Management: (traps)

DECLARE DATA LOSS Trap to let manager know that accounting data is being lost.

DECLARE HIGH WATER Trap to request that manager increase polling interval. (Used when number of flows increases.)

DECLARE FLOOD/FLUSH Trap dumping the flow records currently being monitored by the meter. (Lower priority first?)

Management to meter: (polls and control)

SET HIGH WATER MARK A the meter when to send a trap indicating that the management station should increase the polling interval.

SET FLOOD MARK A how full the table SHOULD be before the meter considers panicking and dumping the contents of the meter to the management station in raw (SNMP OPAQUE) form.

SET FLOW TERMINATION PARAMETERS The meter should have the good sense in situations where lack of resources may cause data loss to purge flow records from its tables which (a) have already been reported and show no activity since the last report (b) oldest flows or (c) flows with the smallest number of unreported packets.

- TIMEOUT The time in seconds since last packet seen (and last report) after which the flow may be terminated.

- MAX LIFETIME Guidelines for the maximum lifetime of a flow. (Not mandatory, but the meter should make an effort at reporting time to purge flows that have had a lifetime greater than this value, even if it results in the instantaneous creation of a new flow with identical parameters.

SET FLOW PRIORITY [REPORTING MASK] (mask is an 8-bit quantity) Tell meter which flows are considered "critical" - i.e., in a crisis situations which flows can least afford to lose data. Reporting mask is set by the RULES TABLE in the SET GRANULARITY operation.

REPORT [REPORTING MASK (0 or default indicates report ALL)] Poll to meter indicating that a normal report of indicated flows should be made (i.e., any flow whose rule has indicated that it has a bit set which is set in the mask.

SET GRANULARITY [RULE TABLE] see RULE TABLE

RULE TABLE: (Editorial comment from the Chair: This is all a very large pie in the sky and not to be sliced seriously yet.)

SEQUENCE OF NUMBERED RULES.

RULE::=
  Field FIELD_DESCRIPTOR.
  [Operator OPERATOR_DESCRIPTOR.]
  Mask. MASK_DESCRIPTOR.
  LIST OF ACTION_PAIRS.

FIELD_DESCRIPTOR ::=
  Length INTEGER. (0 is permitted to indicate lack of interest.)
  CHOICE OF:
    NETWORK ADDRESS. (including arbitrary strings)
    RESULT (VALUE) OF PREVIOUS MASKING OPERATION>.

OPERATOR_DESCRIPTOR ::= The source of much discussion on overhead, complexity, and feasibility. Is anding and testing for equality to the mask good enough or do we need to define a set of allowed operations?

MASK_DESCRIPTOR: A MASK of a length less than or equal to the field. (Otherwise there is no match. 1's set in the mask indicate bits which are of interest. Actually, is defined to be other identical to the field_descriptor. (LENGTH followed by RESULT or NETWORKADDRESS.)

ACTION_PAIR. VALUE or RANGE OF VALUES. If the results of the masking operation fit this value or range of values, perform the following actions.

Choice of:

CONDENSE (FLAGS, FIELD_DESCRIPTOR, [SUBSCRIBER_ID])
EXPAND (FLAGS, FIELD_DESCRIPTOR, {SUBSCRIBER_ID])
IGNORE.

GO TO RULE NUMBER X.

CONDENSE indicates that the flow-record should use the designated FIELD as the source or destination address (or attribute) in the FLOW-ID, along with the designated SUBSCRIBER_ID (also a FIELD_DESCRIPTOR). (Condense implies that all packets parsing to this point will be counted in a single bucket.)

EXPAND is just like condense, except the the FIELD_DESCRIPTOR indicates that the packets which parse to this point should be placed in multiple flows with source or destination address (or attribute) as designated by the the FIELD_DESCRIPTOR.

IGNORE indicates that we don't count this type of packet at all.

USE RULE NUMBER N indicates (theoretically) that the RESULT OF PREVIOUS MASKING OPERATION is set to the result of the FIELD VALUE & (anded with) the mask value, and the nth rule of the RULE TABLE is invoked next. This concept is disturbing because it allows for spaghetti tables that dont make sense. At this point a rule compiler front end becomes necessary...<sigh>

NETWORK_ADDRESS ::= [this should follow reporting format]

Choice of {

[0] IP ADDRESS. (TCP/IP)
[1] NSAP ADDRESS. (OSI) variable length.
[n] X.25 Address (CCITT)
[m] MAC (LLC)
[x] STRING OF OCTETS. (any other arbitrary identifier)}

Notes on Rules

Note that each packet can only by counted within ONE FLOW, so that if all possible flows are added, the number should equal the total number of packets processed.

If there are multiple ways a packet should be processed, the rules should deposit enough information in the flow record (i.e., flow-id) so that the packet can be POST-PROCESSED to be counted in multiple billing categories.

The RESULT field preceding the root of the tree is considered to be a zero length field.

All rule tables must in fact map into non-looping binary trees, or we won't be responsible for the result (To save space sub-trees may be shared by different branches and recursion may be used, as long as it can be shown that no infinite loops can occur Caveat emptor and all that.). When address tests are used (field = address type), recommend performing tests on the interface number first, the link level address second, the network address third, and the attributes (if any are defined later) last. Within an address type, test the source address first and the destination address last.

Attendees

Sean Donelan	sean@dra.com
Martin Dubetz	dubetz@wugate.wustl.edu
Gary Ellis	garye@hpspd.spd.hp.com
Charles Fineberg	fineberg@wums2.wustl.edu
Dave Geurs	dgeurs@mot.com
Keith Hacke	hacke@informatics.wustl.edu
Donald Hirsh	hirsh@meridian.uucp
Cyndi Mills	cmills@bbn.com
Agnes Moran	
Chris Myers	chris@wugate.wustl.edu
Kary Robertson	kr@concord.com.kr
Gregory Ruth	gruth@bbn.com
George Sanderson	sanderson@mdc.com
Jonathan Saperia	saperia@decwrl.enet.dec.com
Anil Singhal	
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Paul Tsuchiya	tsuchiya@bellcore.com
Sudhanshu Verma	verma@hpindbu.cup.hp.com
Steve Witten	

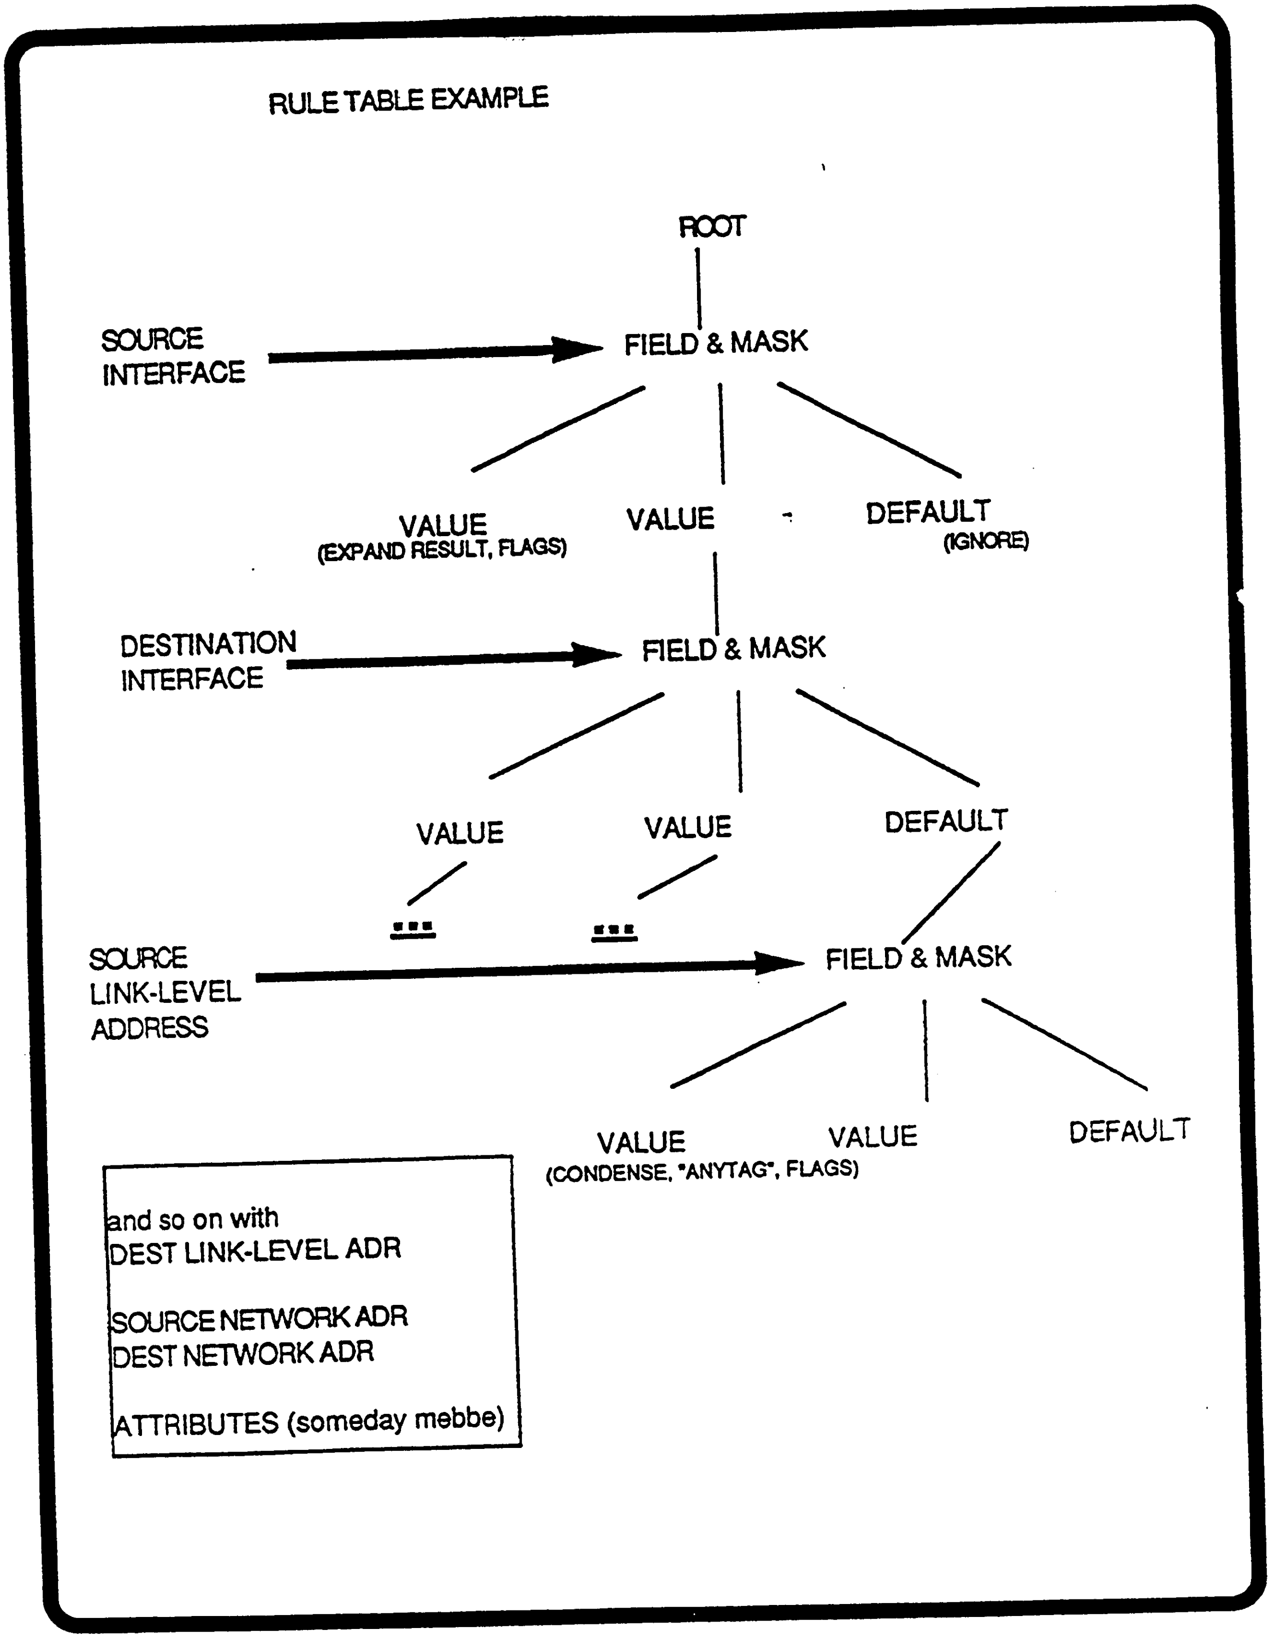
RULE TABLE EXAMPLE
ROOT
SOURCE INTERFACE
FIELD & MASK
VALUE (EXPAND RESULT, FLAGS)
VALUE
DEFAULT (IGNORE)
DESTINATION INTERFACE
FIELD & MASK
VALUE
VALUE
DEFAULT
SOURCE LINK-LEVEL ADDRESS
FIELD & MASK
VALUE (CONDENSE, "ANYTAG", FLAGS)
VALUE
DEFAULT
and so on with
DEST LINK-LEVEL ADR
SOURCE NETWORK ADR
DEST NETWORK ADR
ATTRIBUTES (someday mebbe)

3.3.6 Management Services Interface (msi)

Charter

Chair(s):

Oscar Newkerk, newkerk@decwet.enet.dec.com
Sudhanshu Verma, verma@hpindbu.cup.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 Sudhanshu Verma/HP

MSI Minutes

The MSI Working Group meeting was rescheduled from Tuesday (3/12) afternoon to Monday (3/11) evening. This time clashed with the IETF dinner at the St. Louis zoo. The meeting was attended by only seven people, and was adjourned early for a lack of a quorum.

The issues raised at the last meeting were reiterated. These issues are summarized below.

- On-line MIB database and the need for both GDMO and SNMP MIB definitions.
- Lack of any implementation suggestions or hints in the MSI document. The MSI draft wants the implementors of the MSI API to support features such as translation between SNMP and CMIP and scoping, but does not provide ideas on how to implement this. This has hindered adoption of the API.
- Lack of an SNMP API. Some attendees at the last meeting wanted support of an SNMP-oriented API.

The thorniest issue deals with the issue of translation between the two different SMIs (GDMO and IETF). Based on RFC1109 the IAB has decided that there was no requirement for compatibility between SNMP and OSI network management. This causes the task of the translation between the two SMIs to be done on a case by case basis; it will be difficult, if not impossible to have automated conversion between the two.

The future of the group was also discussed briefly. One option is to work with the IAB and the Network Management Directorate to resolve any pending issues. Another option is to disband the group due to the lack of significant progress. The next Working Group meeting will need to evaluate the situation and make recommendations.

The meeting was adjourned after about 30 minutes due to a lack of a quorum.

Offline

The issue of MIB translation was raised by Sudhanshu Verma with the Chair of the Network Management Area, Chuck Davin. He suggested that there was little that could be done to change RFC 1109 at this point and that we should attempt to work in that framework. The difficulty of defining translation rules to achieve automated translation was reinforced. At the next meeting the Working Group will have to decide its future direction and plans.

Attendees

Steve Bostock	steveb@novell.com
Howard Brown	brown@ctron.com
Shawn Gallagher	gallagher@quiver.enet.dec.com

Ron Mackey
Ron Poppen-Chambers rpc@hpcnd.cnd.hp.com
Sudhanshu Verma verma@hpindbu.cup.hp.com
Mark Wood

3.3.7 OSI Internet Management (oim)

Charter

Chair(s):

Lee LaBarre, cel@mbunix.mitre.org
Brian Handspicker, bd@vines.enet.dec.com

Mailing Lists:

General Discussion: oim@mbunix.mitre.org
To Subscribe: oim-request@mbunix.mitre.org

Description of Working Group:

This Working Group 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.

3.3.8 Remote LAN Monitoring (rlanmib)

Charter

Chair(s):

Mike Erlinger, `mike@mti.com`

Mailing Lists:

General Discussion: `rlanmib@mti.com`
To Subscribe: `rlanmib-request@mti.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:

Done	Mailing list discussion of Charter and collection of concerns.
Done	Discussion and final approval of Charter; discussion and agreement on models and terminology. Make writing assignments.
Done	Discussion of the first draft document. Begin work on additional drafts if needed.
Mar 1991	Review latest draft of the first document and if OK give to IESG for publication as an RFC.

CURRENT MEETING REPORT

Reported by Michael Erlinger/Micro Technology

Remote LAN Monitoring Minutes

Three separate meetings were held with the primary Agenda to review the RLANMIB MIB proposed by Steve Waldbusser. The MIB had been distributed to the mailing list and copies were available at the meeting. The driving focus of the current MIB is to quickly get a consensus on an RLANMIB that can act as a standard. For this reason, various issues have been moved to future MIBs. In general the MIB document should have more verbiage describing the MIB and the general philosophy that was followed.

Memory management and table size issues were discussed at length. The only consensus reached is that memory management is a problem and that the various probes will find their own way to control memory.

A philosophic question was raised and not debated: What is the difference between a monitor and an analyzer? This needs to be discussed more to better decide on the RLANMIB.

During the discussions about multiple managers and table ownership, the point was made that the probability of multiple manager collisions was in fact quite high, since access to probe tables is often the result of network problems (during which more than one manager may rush to fix). MIB development needs to recognize this point. It was decided that a RLANMIB meeting should be held prior to the next IETF. The date of this meeting will be decided after a new version of the MIB document is made available to the mailing list. The Chair will be responsible for choosing the date and location.

A few general points were discussed as foundation principles for the RLANMIB:

- Probes will be used simultaneously by more than one network management station.
- Probe resources will be a constant concern, a method must be found that would allow a probe to determine which dynamic tables, particularly those associated with an NMS, can be deleted.
- Accepting the simultaneous use of the probe, the MIB should insure the isolated use by each NMS.
- Accepting the simultaneous use of the probe, the MIB should allow for the sharing of use by each NMS.

MIB Review

Etherstats Table

Various entries are the same as other MIBs, (ethernet), while other entries are new. Two justifications for this approach:

1. Probes have the primary task of monitoring so the additional resources should not be a concern;

2. Probes operate in promiscuous mode, so they will produce different values.

MIB should spell out whether good and/or bad packets are included in a count. In general this information should be added to all counter descriptions. MIB should spell out that: All counts exclude framing - start with destination address and continue through CRC. In particular, all packets are included in each bucket because segment utilization includes both good and bad packets.

Etherstats Counters

```
          ‘64                   64--1518                 ˜1518
        |--------------------------------------------------------------
CRC     | collision               crc/align                jabber
error   | fragments
        |--------------------------------------------------------------
NONE    | runt                    good                     oversize
```

It was noted that the etherSTatsPkts64Octets counter was missing – to be added in next version.

Inclusive or exclusive will be added to text describing various packet counters.

Etherhistory Table

Circular rollover: when the N buckets are full you continue to have only N buckets, loosing the oldest bucket.

Interval change semantics: It is viewed that a change in interval is the same as deleting the current control entry and starting a new one, i.e., the existing N buckets are lost and new N buckets with the current interval are allowed to exist in the system (actual allocation of buckets is an agent task).

Change # of buckets semantics: Changing the number of buckets should not invalidate the current buckets. This will be explained in the document. In particular, changing to a greater number of buckets, just adds more buckets to that history sequence. Reducing the number of buckets deletes the oldest buckets until the required number are left.

What about time stamping the bucket contents? Adding an end time to the bucket has little meaning because granularity is probably 1 sec and is thus not very meaningful. Buckets are not real-time. Finally, could use the start time of the next bucket as the end time of the current bucket.

Discussion of starting a table entry: The entry starts when the VALID is set. Valid could be set in the same PDU as all the other entries because a set is viewed as an atomic operation.

How to determine if probe lost data: Use dropevents to determine if probe lost some events.

Utilization will be changed from tenths of a percent to hundredths of a percent.

Utilization discussion: Because everyone determines utilization differently (some use various hardware tests), it was decided that the utilization value is a standard way of presenting a non-standard value.

A request was made for max available history buckets counter (etherHistory 3??). Someone said this is necessary in all dynamic tables and quite useful for the management station user interface

Ether Host Tables

The etherhostorder table will be ordered by time of 1st transmission – still 1 to N. Much discussion and much debate about the problem of deleting stations from the table and still maintaining the ordering. This is an open issue which must be explained in the document.

The host table ordered by natural index is being used to serve two purposes: fast download of the whole table and new station detection. The first requires a contiguous index space (necessitating renumbering) and the second requires monotonically increasing indices. The resolution was to create two tables instead of one (although Steve said he would try to figure out a way to shrink them back into one table).

Table deletion: It was decided that most tables need a deletion capability and that the MAC address is the most secure way to do deletion. Other indices may actually change.

After much debate about the TOP N table, it was decided that there are three options:

1. Leave the table as it currently is;
2. Nuke the whole idea;
3. Expand the table to a series of tables – a control table that describes each of the actual top N tables.

Some discussion about probe reaction when a table that is already "valid" is set to "valid". It was decided that the proper agent action is "no-op".

Ether Traffic Matrix Tables

Change "etherSDTableSize" to "etherMatrixTableSize"

The Filter table again raised the issue of NMS control of specific tables and the Probe/Agent's ability to garbage collect.

The idea of an X.Y index for each dynamic NMS related table was discussed. A central table would exist in which each NMS specified its own unique X value. The NMS would also specify the time in sections for which X related tables should be maintained by the Agent. If the time decrements to 0, the Agent can reclaim all tables and table entries related to the NMS. The NMS can periodically restart the countdown clock. Thus, an NMS knowing its own unique ID, can keep all its tables active (since it knows the time value entered into the station), the NMS can force deletion of all its tables by entering 0 into the time field, and yet the Agent can delete tables related to a particular NMS that is no longer active. Also, if this table includes the IP address or some other know NMS address, all NMSs can determine what other NMSs are using dynamic tables in the agent.

Buffer Control Table

Steve will add a variable to the buffer control table that holds the max available entries in the capture table.

There was some discussion about how an agent would treat a set request in which either (or both) bufferControlCaptureSliceSize or bufferControlUsedBuffers were zero. Does this constitute a space reservation? Can the agent return BadValue? No resolution of this question.

All state variables in control tables will have an Invalid state added (enfferControlState == stopWhenFull) implies that any filters which are supposed to "turnOn" that buffer, will not, once the buffer has reached the full state.

A variable should be added to the bufferControlTable containing the value of sysUpTime when the buffer was first turned on.

The syntax of "captureBufferPacketTime" will be changed from TimeTicks to an integer containing the number of milliseconds since the buffer was turned on.

Steve stated that the intent of the log table was to keep the most recent events (once it started wrapping).

There was some discussion on numbering traps in the global environment. Steve will give it some more thought.

We need to add a notificationIndex to the logTable.

Notification Table

A minimal time was spent discussing the Notification Table. Traps were referenced to the Notification Table, but not discussed in detail.

Off Line

Overhead associated with updating the etherhistory table.

Steve will write up a mail message that will explain his approach to fast table dumping and the type of performance that he obtained.

Topics for Later Discussion

A table describing thresholds will be added at a later time, hopefully in the next version of the MIB.

Deferred Topics

Various topics are being deferred to RLANMIB 2 in the interest of expediently getting a RLANMIB 1 into test and evaluation.

Peaks: Peaks are difficult especially in handling sliding time windows. If discrete time is used, then it is possible to miss peaks. In determining which type of peaks will be captured, e.g., utilization, broadcasts, etc., it was realized that peaks could double the size of the history table. Peaks should be time tagged, not just captured in the history table. Peaks really fall into the threshold area.

The concept of protocol bitmasks for each station and protocol percentages for the segment were discussed at length. The consensus was to let this area go until RLANMIB 2. Questions that seemed open: protocols to be included in the bitmask; how far down the stack protocol counting occurs; and general utilization of this feature.

Attendees

William Anderson	wda@mitre.org
William Barns	barns@gateway.mitre.org
Steve Bostock	steveb@novell.com
Kurt Dobbins	dobbins@ctron.com
Bill Durham	durham@MDC.COM
Gary Ellis	garye@hpspd.spd.hp.com
Fred Engel	engel@concord.com
Mike Erlinger	mike@mti.com
Bill Fardy	fardy@ctron.com
Martin Gray	mg@spider.co.uk
Mike Janson	mjanson@mot.com
Kenneth Key	key@cs.utk.gdy
Cheryl Krupczak	clefor@secola.columbia.ncr.com
Donna McMaster	dmcmaster@synoptics.com
Lynn Monsanto	monsanto@eng.sun.com
David Perkins	dave_perkins@3com.com
Ron Poppen-Chambers	rpc@hpcnd.cnd.hp.com
Rehmi Post	rehmi@ftp.com
Ron Roberts	roberts@jessica.stanford.edu
Kary Robertson	kr@concord.com.kr
Jonathan Saperia	saperia@decwrl.enet.dec.com

Greg Satz	satz@cisco.com
Mark Schaefer	schaefer@davidsys.com
Dror Shindelman	pbrenner@sparta.spartacus.com
Anil Singhal	
Sudhanshu Verma	verma@hpindbu.cup.hp.com
Steven Waldbusser	steven.waldbusser@andrew.cmu.edu
Steve Witten	
Wing Fai Wong	wfwong@malta.sbi.com

3.3.9 SNMP Security (snmpsec)

Charter

Chair(s):

James Galvin, galvin@tis.com
Keith McCloghrie, kzm@hls.com

Mailing Lists:

General Discussion: snmp-sec-dev@tis.com
To Subscribe: snmp-sec-dev-request@tis.com

Description of Working Group:

The SNMP Security Working Group is chartered to determine the set of security services needed by the SNMP. The specification of those services, the supporting mechanisms, and the adjunct infrastructure will become an enhancement to the SNMP and eventually an Internet standard.

The specification must not alter the fundamental SNMP network management philosophy and must not entail changes to existing SNMP standards or framework.

Goals and Milestones:

Done	Publish internet-draft specifications.
Jul 1991	Submit specification to IESG for consideration as a Proposed Standard.
Dec 1991	Submit specification to IESG for consideration as a Draft Standard.
Ongoing	Submit specification to IESG for consideration as a Standard.

CURRENT MEETING REPORT

Reported by James Galvin/TIS and Keith McCloghrie/Hughes

SNMPSEC Minutes

The Working Group met for one evening to discuss the latest revision of the documents. There are currently three documents:

1. SNMP Administrative Model - this document specifies a framework within which the protocols specified above funtion.

2. SNMP Security Protocols - this document is a combination of what was previously the "Authentication and Privacy" document and the "Administration" document. It completely specifies two enhancements to SNMP to support integrity and authentication, and integrity, authentication and privacy.

3. SNMP Party MIB - this document specifies a set of experimental MIB objects that may be used to support the SNMP administrative model, including the SNMP security protocols specified above.

All three documents require some editorial changes, after which they will be submitted to the Internet-Drafts directory. Insofar as the second document represents changes to the SNMP that are not strictly part of the Working Group Charter, the proper process for further progressing of the document is a matter for IESG consideration. Since the protocol document depends on the model document, it may not progress to a proposed standard until the status of the model document is resolved.

The set of slides used to present the administrative model and the changes to the security protocols are included below.

Attendees

Karl Auerbach	karl@eng.sun.com
Jordan Becker	becker@nis.ans.net
Steve Bostock	steveb@novell.com
David Bridgham	dab@asylum.sf.ca.us
Howard Brown	brown@ctron.com
Theodore Brunner	tob@thumper.bellcore.com
Christopher Bucci	bucci@pluto.dss.com
Jeffrey Case	case@cs.utk.edu
James Davin	jrd@ptt.lcs.mit.edu
Kurt Dobbins	dobbins@ctron.com
Fred Engel	engel@concord.com
Shari Galitzer	shari@gateway.mitre.org
Shawn Gallagher	gallagher@quiver.enet.dec.com

James Galvin	galvin@tis.com
Joseph Golio	golio@cray.com
Martin Gray	mg@spider.co.uk
Jeremy Greene	greene@coral.com
Frank Kastenholz	kasten@asherah.clearpoint.com
Kenneth Key	key@cs.utk.gdy
Jim Kinder	jdk@fibercom.com
Stev Knowles	stev@ftp.com
John Linn	linn@zendia.enet.dec.com
Keith McCloghrie	kzm@hls.com - his preferred address
Lynn Monsanto	monsanto@eng.sun.com
Michael Patton	map@lcs.mit.edu
David Perkins	dave_perkins@3com.com
Jonathan Saperia	saperia@decwrl.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
Jim Sheridan	jsherida@ibm.com
Dror Shindelman	pbrenner@sparta.spartacus.com
Dean Throop	throop@dg-rtp.dg.com
Steven Waldbusser	steven.waldbusser@andrew.cmu.edu
Wengyik Yeong	yeongw@psi.com
Jeff Young	jsy@cray.com

SNMP Security

PROBLEMS:

- limit of 100 messages per second
- fuzzy notion of source / target of request
- reality of our claim of future "extension" to asymmetric cryptography
- relationship to other aspects of network management framework i.e., access control, proxy, MIB views
- readability and understandability of documents

PROPOSED SOLUTIONS:

- one change to Security Protocols – the "nonce"
- specific and separate identification of source and destination in Message
- re-organization of text / explanations
- elaboration of SNMP Administrative Model
- "practice makes perfect"

Source and Destination

- authentication based on source
- privacy based on destination
- access control based on pairing of source and destination
- proxy based on destination
- MIB view based on destination

SNMP Party

- An "execution context" within an SNMP protocol entity
- A restricted subset of operations, i.e., a subset of what can be done by the SNMP protocol entity

SNMP Party (continued)

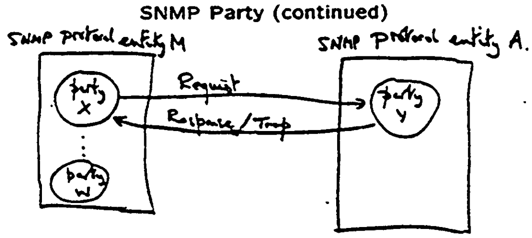

- **Symmetric Cryptography:** both M & A must know party "secrets" and other parameters of both X & Y
- **Asymmetric Cryptography:** both M & A must know some information about X & Y, but only M knows X's secrets, and only A knows Y's secrets (possible hook for the future)

SNMP Party (continued)

- Identified by an OBJECT IDENTIFIER
- Has authentication parameters used when this party is a source
- Has privacy parameters used when this party is a destination
- "Lives" at a transport address of some transport stack, e.g., SNMP/UDP/IP
- A single MIB view
- Local or Proxy

Party Naming

- Unique OBJECT IDENTIFIER for each party
- Initial sets of party identifiers defined, by convention, 6 for each IP address

{initial1157PartyId a b c d 1 }	noAuth	noPriv
{initial1157PartyId a b c d 2 }	noAuth	noPriv
{initial1157PartyId a b c d 3 }	md4Auth	noPriv
{initial1157PartyId a b c d 4 }	md4Auth	noPriv
{initial1157PartyId a b c d 5 }	md4Auth	desPriv
{initial1157PartyId a b c d 6 }	md4Auth	desPriv

7 SNMP Security

Party Authentication Information

- Algorithm: noAuth, md4Auth
- Keys: partyAuthPrivate, partyAuthPublic, partyAdmin
- Clocks: partyAuthClock, partyAuthLastMsg, partyAuthNonce, partyAuthLifetime, partyAdminClock

8 SNMP Security

Party Privacy Information

- Algorithm: noPriv, desPriv
- Keys: partyPrivPrivate, partyPrivPublic

9 SNMP Security

MIB Views

- one-to-one with SNMP party
- named by party identifier
- set of mutually disjoint view subtrees

10 SNMP Security

View Subtree

- Named by a node in the MIB tree
- Contains all possible object instances within MIB tree

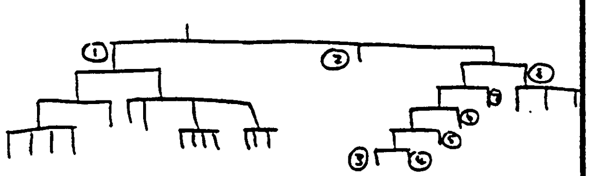

SNMP Security

Access Control

- A source party
- A destination (target) party
- A set of "Management Communication Classes" i.e., a set of PDU types

Access is permitted for a particular source a particular target to process a set of request types

12 SNMP Security

Proxy

- an OBJECT IDENTIFIER
  - either **noProxy** – operations performed locally
  - or **pointer to another party** – other party information about transport stack and transport address of proxied agent

13 SNMP Security

Message Formats

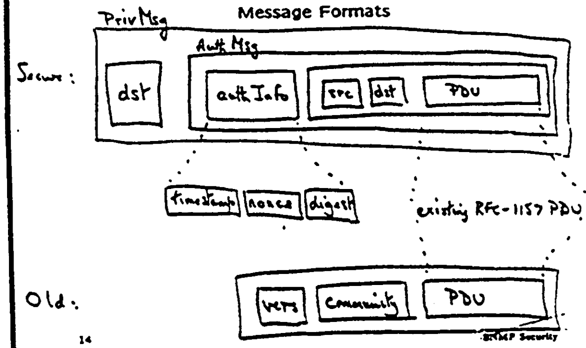

14 SNMP Security

Compatibility

- As compatible as reasonably possible
- Same PDU – different Message wrapper
- No surprises here – forecast by RFC 1157, RFC 1098, and RFC 1067 dating back to May, 1988
- Agents implement one (old style or secure) – preferably secure
- Managers implement both – at least two ways to do this

15 SNMP Security

And, Finally

- After 2 years, at least a dozen documents, > 10 M bytes of archives, and loss of many trees and sleep

WE ARE DONE. (whew!)

16 SNMP Security

snmp-sec-dev@tis.com

TO GET ADDED TO THE LIST

snmp-sec-dev-request@tis.com

3.3.10 Simple Network Management Protocol (snmp)

Charter

Chair(s):
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:

Oversee development of SNMP-related activity, especially the Internet-standard SMI and MIB. This Working Group is ultimately responsible for providing workable solutions to the problems of network management for the Internet community.

Goals and Milestones:

Aug 1990 Finish SNMP Authorization draft.

Ongoing Coordinate the development of various experimental MIBs.

CURRENT MEETING REPORT

Reported by James Davin/MIT

SNMP Minutes

The SNMP Working Group of the IETF met on Monday morning, March 11, 1991. J. Davin acted as Chair in lieu of Marshall Rose, who was unable to attend.

The Working Group considered two documents:

- The MIB for Ethernet-like Interfaces `<draft-ietf-snmp-ethernetmib-03.txt>` and
- The SMDS Interface Protocol MIB `<draft-ietf-snmp-smdsipmib-00.txt>`

A brief presentation of the Ethernet MIB was made by J. Davin to refamiliarize Working Group members with the current status of this longstanding effort. There followed a section-by-section discussion of the document. The Working Group decided on several minor changes:

- Two typographical errors in enumerated type definitions were corrected.
- Text was added to the document to guide users of the MIB in how to interpret error counts when an agent may be unable to accurately report all instances of certain error conditions.
- The set of names of Ethernet-like chipsets (for use with the Interface Extensions MIB) was enlarged.

With these changes, the Working Group recommended that the document be considered by the IESG for Proposed Standard status.

In the second half of the meeting, Kaj Tesink introduced the group to the SMDS Interface Protocol MIB by giving an overview presentation on the SIP itself and the structure of the proposed MIB. A discussion of the document followed. Some concerns were raised about the relationship of this MIB to the interfaces group of MIB II, about the policy questions involved in the choice of objects, about syntactic and naming conventions used in the MIB, and about the clarity of some of the counter definitions. At the end of the discussion, the Working Group decided to consider this MIB further in subsequent meetings.

Attendees

William Anderson	wda@mitre.org
Karl Auerbach	karl@eng.sun.com
Ballard Bare	bare@hprnd.rose.hp.com
William Barns	barns@gateway.mitre.org
Steve Bostock	steveb@novell.com
Andreas Bovopoulos	andreas@patti.wustl.edu
Richard Bowles	bowles@stsci.edu
David Bridgham	dab@asylum.sf.ca.us
Theodore Brunner	tob@thumper.bellcore.com
Christopher Bucci	bucci@pluto.dss.com
Charles Carvalho	charles@sage.acc.com
Jeffrey Case	case@cs.utk.edu
Don Coolidge	coolidge@speaker.wpd.sgi.com
Tracy Cox	tacox@sabre.bellcore.com
James Davin	jrd@ptt.lcs.mit.edu
Kurt Dobbins	dobbins@ctron.com
Nadya El-Afandi	nadya@network.com
Gary Ellis	garye@hpspd.spd.hp.com
Mike Erlinger	mike@mti.com
Bill Fardy	fardy@ctron.com
Charles Fineberg	fineberg@wums2.wustl.edu
Karen Frisa	karen.frisa@andrew.cmu.edu
Shawn Gallagher	gallagher@quiver.enet.dec.com
James Galvin	galvin@tis.com
Joseph Golio	golio@cray.com
Keith Hacke	hacke@informatics.wustl.edu
Patrick Heisinger	heisinger@mdcguy.mdc.com
Mike Janson	mjanson@mot.com
Frank Kastenholz	kasten@asherah.clearpoint.com
Kenneth Key	key@cs.utk.gdy
Cheryl Krupczak	clefor@secola.columbia.ncr.com
Nik Langrind	nik@shiva.com
John LoVerso	loverso@westford.ccur.com
Ron Mackey	
Keith McCloghrie	kzm@hls.com - his preferred address
Donna McMaster	dmcmaster@synoptics.com
Carol Melowitz	melowitz@mdcgwy.mdc.com
Jay Melvin	infopath@well.sf.ca.us
Linda Melvin	infopath@well.sf.ca.us
Donald Merritt	don@brl.mil
Lynn Monsanto	monsanto@eng.sun.com
Bill Norton	wbn@merit.edu
Robert Peglar	robp@anubis.network.com

David Perkins	dave_perkins@3com.com
Ron Poppen-Chambers	rpc@hpcnd.cnd.hp.com
Jonathan Saperia	saperia@decwrl.enet.dec.com
Mark Schaefer	schaefer@davidsys.com
Dror Shindelman	pbrenner@sparta.spartacus.com
Anil Singhal	
Evan Solley	solley@applelink.apple.com
Bob Stewart	rlstewart@eng.xyplex.com
Kaj Tesink	kaj@nvuxr.cc.bellcore.com
Dean Throop	throop@dg-rtp.dg.com
Maurice Turcotte	dnmrt@rm1.uucp
Sudhanshu Verma	verma@hpindbu.cup.hp.com
Steven Waldbusser	steven.waldbusser@andrew.cmu.edu
Steve Witten	
Wing Fai Wong	wfwong@malta.sbi.com
Mark Wood	
Wengyik Yeong	yeongw@psi.com
Jeff Young	jsy@cray.com
Osmund deSouza	desouza@osdpc.ho.att.com

3.3.11 X.25 Management Information Base (x25mib)

Charter

Chair(s):

Dean Throop, throop@dg-rtp.dg.com

Mailing Lists:

General Discussion: x25mib@dg-rtp.dg.com
To Subscribe: x25mib-request@dg-rtp.dg.com

Description of Working Group:

This working group will produce a set of three documents that describe the Management Information Base for X.25. The first document will specify the objects for the X.25 Link Layer. The second document will specify the objects for the X.25 Packet Layer. The third document will specify the objects for managing IP over X.25. The working group need not consider the Physical Layer because the "Definition of Managed Objects for RS-232-like Hardware Devices" already defines sufficient objects for the Physical Layer of a traditional X.25 stack. Any changes needed at the Physical Layer will be addressed as part of that activity.

The X.25 object definitions will be based on ISO documents 7776 and 8208 however nothing should preclude their use on other similar or interoperable protocols (i.e. implementations based on CCITT specifications).

The objects in the Link and Packet Layer documents, along with the RS-232-like document, should work together to define the objects necessary to manage a traditional X.25 stack. These objects will be independent of any client using the X.25 service. Both of these documents assume the interface table as defined in MIB-II contains entries for the Link and Packet Layer interfaces. Thus these documents will define tables of media specific objects which will have a one to one mapping with interfaces of ifType ddn-x25, rfc877-x25, or lapb. The objects for the IP to X.25 convergence functions will be defined analogously with the ipNetToMedia objects in MIB II.

The working group will endeavor to make each layer independent from other layers. The Link Layer will be independent of any Packet Layer protocol above it and should be capable of managing an ISO 7776 (or similar) Link Layer provider serving any client. Likewise the X.25 Packet Layer objects should be independent of the Link Layer below it and should be capable of managing an ISO 8208 (or similar) Packet Layer serving any client.

The working group will also produce a third document specifying the objects for managing IP traffic over X.25. These objects will reside in their own table but will be associated with the X.25 interfaces used by IP. These objects will not address policy decisions or other implementation specific operations associated with X.25 connection management decisions except as explicitly described in existing standards. These objects will manage the packet flow between IP and the X.25 Packet Layer specifically including

observation of packet routing and diagnosis of error conditions. Progress on the Link and Packet Layer documents will not depend on progress of the IP over X.25 document. The IP over X.25 document will proceed on a time available basis after work on the Link and Packet Layer documents and as such the Link and Packet Layers may be completed before the IP over X.25 work.

All documents produced will be for use by SNMP and will be consistent with other SNMP objects, conventions, and definitions (such as Concise MIB format). To the extent feasible, the object definitions will be consistent with other network management definitions. In particular ISO/IEC CD 10733 will be considered when defining the objects for the X.25 Packet Layer.

Goals and Milestones:

Apr 1991	Distribute first draft of documents and discuss via E-mail.
Aug 1991	Working group meeting as part of IETF to review documents.
Sep 1991	Distribute updated documents for more E-mail discussion.
Nov 1991	Review all documents at IETF meeting. Hopefully recommend advancement with specified editing changes.
Jan 1992	Documents available with specified changes incorporated.

3.4 OSI Integration Area

Director(s):

- Ross Callon: callon@bigfut.enet.dec.com
- Rob Hagens : hagens@cs.wisc.edu

Area Summary reported by Rob Hagens/UWISC

At the St. Louis IETF meeting, three Working Groups from the OSI Integration Area met. A fourth group, OSI-DS, met in February.

OSI-ODA

The Office Document Architecture Working Group held its inaugural meeting at the St. Louis IETF. This group will liaise with the SMTP extensions Working Group, the X400 Working Group, and the Netfax Working Group. The group has agreed to set up a small pilot by the next IETF. The pilot will initially be based upon the Slate+ODA software (running on a sparc station) and WordPerfect software (running on a PC). The ODA group will produce a document on the use of SMTP and X.400 to carry ODA documents; this paper will be available in May. The group will investigate other available products to add to the pilot. The ODA group invites other Working Groups who need to exchange revisable text/diagrams/bitmaps to join the ODA pilot project.

OSI-X400OPS

The X.400 Operations Working Group also held their inaugural meeting at the St. Louis IETF. Participants at this meeting included vendors as well as X.400 planning staff who serve user communities. The group made concrete progress in several areas: agreement on preliminary X.400 address registration mechanisms; agreement on basic service documentation with mechanisms to insure full end-user connectivity, routing information exchange, and documentation of reachability; agreement on the need for one common rule for X.400/RFC 822 address mapping for the portion of the Internet within the US, and a draft decision on a mapping rule. The result of this meeting is the outline of Internet PRMD requirements. A draft version of the document produced from this outline will be available in June.

OSI-NOOP

The Network OSI Operational (NOOP) Working Group also held their inaugural meeting at the St. Louis IETF. The group discussed their Charter. They added OSI tutorials relevant to the Working Group's activities as an additional publication category. The group discussed sample routing plans and received volunteers to document various plans. These will be used as the basis of future discussion and as examples for others who are contemplating NSAP addressing and routing planning. The group discussed various categories of debugging tools that are available or need to be written.

OSI-DS

The OSI-DS group met in February. The Minutes of this meeting have been included in these Proceedings. The vast number of topics covered at this meeting is too long to reproduce in this area overview. The list includes liaison reports, technical discussions, pilot project status reports, and document review. The interested reader should consult the attached Minutes for more detailed information.

3.4.1 Assignment of OSI NSAP Addresses (osinsap)

Charter

Chair(s):

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:

Done	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 User Guide" and the FNC OSI Planning Group document.

3.4.2 Network OSI Operations (noop)

Charter

Chair(s):

Susan Hares, skh@merit.edu

Mailing Lists:

General Discussion: noop@merit.edu
To Subscribe: noop-request@merit.edu

Description of Working Group:

Focus on the Operational issues of deploying OSI Network Layer in the Internet.

Develop OSI Routing Plans

Improve Management of the OSI Network Layer

Goals and Milestones:

TBD	Collect OSI Routing and Addressing plans into a Repository
TBD	Provide a forum to discuss these OSI Routing plans by email or in group discussions
TBD	Collect a list of OSI Network Utilities available in the public domain and from vendors. This list will be passed over to the NOC tools group effort for joint publication.
TBD	Collect list of OSI Network Layer NOC tools and publish a list.
TBD	Collect Methods of OSI Network Layer Debugging and write a document describing these methods

CURRENT MEETING REPORT

Reported by Sue Hares/Merit and Richard Colella/NIST

NOOP Minutes

The first meeting of the Network OSI Operational (NOOP) Working Group was opened by the Chair, Sue Hares (Merit). A minute-taker was selected, introductions were made, and the Agenda was presented and approved as presented.

Sue presented the NOOP Working Group proposed Charter for discussion. It was agreed that an additional category be added to the areas for publication, that of tutorials on OSI relevant to the Working Group's activities. Several of these were identified immediately:

- DIS 10589 - Ross Callon (DEC) agreed to find or write a tutorial on DIS 10589, the IS-IS Intra-domain Routing protocol;
- IDRP - Dave Katz (Merit) will write a tutorial on the Inter-domain Routing Protocol, currently under consideration in ANSI and ISO; and,
- CLNP and ES-IS - Rob Hagens (U. Wisc) will allow us to post two tutorial articles he's written, and previously published in Connexions on CLNP and ES-IS.

Tutorial on OSI NSAP Guidelines

As an introduction to the discussion of routing plans, Ross presented a tutorial on the Internet Draft document, "Guidelines for OSI NSAP Allocation in the Internet" (Colella, Gardner, Callon). The slides will be included in the Proceedings.

Five Questions about NSAPs

Sue addressed the five most-often asked questions about NSAPs:

1. How does one get NSAPs outside the US?

   The authority for NSAP allocation under the Data Country Code (DCC) is assigned by ISO to the ISO Member Body (MB) from each country. In the US, for example, this is ANSI; in the UK it is the British Standards Institute. Each MB is responsible for administering (or delegating the adminis- tration of) NSAPs under its DCC.

   Alternatively, one could approach an organization that has obtained an International Code Designator (ICD), which are not country-based.

2. What happens when you change regionals?

   If your NSAP addresses are taken from a regional's address space and you change regionals, your addresses will need to change. This is a consequence of using hierarchical addressing.

   Auto-configuration of ES NSAPs will help this, but there is no standard on this to date (but see below).

X.500 must be capable of supporting this transition. Since the primary source of addressing information for communicants will be acquired through X.500, X.500 must have mechanisms to support wholesale changes of NSAPs.

3. What is the state of auto-configuration of NSAPs?

   Today, some vendors use the ES-IS HELLO PDU to construct ES NSAPs. The ES listens for a HELLO when it is booted, and uses the area address of the IS for its area address. The ES uses its own IEEE MAC address in the system ID field of the NSAP.

   There is a new work item in ANSI X3S3.3 for support of address transition. The proposed mechanism is to use IS-IS flooding and ES-IS to distribute new NSAP information within an area. Note that this does not obviate the need for some support from X.500 as discussed in the previous question.

4. Quality of Service (QoS) – what about it?

   QoS is also known by other names, such as Type of Service (ToS) and policy. This is considered a research issue (and is for further study by someone else).

5. How does a company transition between two different OSI NSAPs?

   It was decided that this question was simply a way of asking question 2.

Sample Routing Plans

There was general agreement on several points:

- Regionals and Backbones should each obtain an AA from GSA or an ORG ID from ANSI and allocate addresses to their clients based on these (Richard agreed to distribute information on how to obtain an AA to the Working Group mailing list).
- Multiply-homed clients (e.g., campuses) should use addresses taken from the address of their primary point of attachment.
- For administrative reasons, a regional *could* further subdivide the RD field. For example, MIDNET clients are organized along state boundaries. The first nibble of the RD field could be used as a state identifier, with the other 12 bits managed by the state. It was decided that this is not a good idea:
  - Each additional subdivision wastes address space, so this should only be done when absolutely necessary,
  - Does not provide a significant added value to the state (new RDs are not needed that often),
  - It may cause confusion because current administrative fields in the NSAP DSP (i.e., AA, RD, and Area) are being used synonymously with topological routing

structures, and these bits are *not*, and,

- It's really just as easy for the regional to manage a flat RD space.

Seven people agreed to write routing plans and submit them for distribution to the Working Group. These plans can be used as the basis for discussion of issues that arise and as examples for others who are contemplating developing NSAP addressing and routing plans. No specific dates were set, but the following lists those who agreed to contribute a plan:

CICNET – Linda Winkler (Argonne National Lab)
ESNET – Tony Hain (Lawrence Livermore National Lab)
MIDNET – Dale Finkelson (MIDNET)
Mitre – Walt Lazear (Mitre)
NSFNET – Dave Katz (Merit)
OARNET – Kannan Varadhan (Ohio Academic Resource Network)
Westnet – Carol Ward (University of Colorado)

Tools

As in the operational IP networks, there is a need for OSI tools to support network operation and debugging. Sue broke this up into a number of separate areas.

- Utilities

  Various OSI-based utilities are needed. Some are available and some are not. These need to be identified and implemented. The slides identified a number of utilities. Note that network management is not of immediate concern to the Working Group, but CMIP-based NM is anticipated.

  [Editor's note: The text of the slides will be sent to the mailing list. The basic utilities needed were:

  ISO versions of ping (ISO echo RFC 1139),
  ISO version of traceroute (using ISO ping as packet sent)
  Method to display ISO routes
  Method to display ES-IS cache

  Making NM Monitors

  Utilities such as ISO ping can be combined to create Network Management Monitors. Sue Hares will send a write up of a simple network management monitor to the mailing list.

- NOC Tools 2

  This Working Group will collect information on OSI tools. A document point to OSI utilities in either the NOC tools RFC will be written. Please send information about OSI utilties to noop@merit.edu.

- Router and Host Survey Documents

  It would be useful to have a pair of documents that survey OSI software availability, one each for hosts and routers. Emphasis is on those tools that are needed to run an operational network. Everyone is asked to contribute material to this effort. Again send information to noop@merit.edu.

Action Items

This section contains a summary of action items from the St. Louis meeting.

DIS 10589 tutorial – Ross Callon
IDRP tutorial – Dave Katz
CLNP and ES-IS tutorials – Rob Hagens
Info on obtaining a GOSIP AA – Richard Colella
CICNET routing plan – Linda Winkler
ESNET routing plan – Tony Hain
MIDNET routing plan – Dale Finkelson
Mitre routing plan – Walt Lazear
NSFNET routing plan – Dave Katz
OARNET routing plan – Kannan Varadhan
Westnet routing plan – Carol Ward
Router and Host survey information – all
Write-up on simple ISO pingky monitor - Sue Hares

Attendees

Richard Bowles	bowles@stsci.edu
Ross Callon	callon@bigfut.enet.dec.com
Lida Carrier	lida@apple.com
Eric Carroll	eric@utcs.utoronto.ca
Martina Chan	mchan@mot.com
Richard Colella	colella@osi3.ncsl.nist.gov
Tom Easterday	tom@cic.net
Dennis Ferguson	dennis@canet.ca
Dale Finkelson	dmf@westie.unl.edu
Daniel Friedman	danfriedman@bbn.com
Vince Fuller	vaf@Standford.EDU
Demi Getschko	"DEMI@FPSP.HEPNET"
Fred Gray	fred@homer.msfc.nasa.gov
Tony Hain	alh@eagle.es.net
Susan Hares	skh@merit.edu
Kathleen Huber	khuber@bbn.com
Mike Janson	mjanson@mot.com

Ruth Lang	rlang@nisc.sri.com
Walter Lazear	lazear@gateway.mitre.org
E. Paul Love	loveep@sdsc.edu
Bill Manning	bmanning@houston.sc.ti.com
Linda Melvin	infopath@well.sf.ca.us
David Miller	dtm@ulana.mitre.org
Dennis Morris	morrisd@imo-uvax.dca.mil
Donald Morris	morris@ucar.edu
John Reinart	reinart@cray.com
Michael Tharenos	tharenos@jessica.stanford.edu
Kannan Varadhan	kannan@oar.net
Ross Veach	rrv@uiuc.edu
Shujiuan Wang	swang@ibm.com
Carol Ward	cward@spot.colorado.edu
Chris Weider	clw@merit.edu
Linda Winkler	lwinkler@anl.gov
Cathy Wittbrodt	cjw@nersc.gov
Wengyik Yeong	yeongw@psi.com

NSAP Address Guidelines

Ross Callon
3/91

NSAPS

WHAT IS THE ISSUE ?

BACKGROUND ON OSI ROUTING

ASSUMPTIONS / OBSERVATIONS

TECHNICAL DETAILS

RECOMMENDATION

ISSUE

- ISO 8473 (CLNP) is OSI EQUIVALENT TO IP
  - LONGER, FLEXIBLE ADDR.
  - INTERNATIONALIZED
- USE WITH: ES-IS / IS-IS ROUTING / INTERDOMAIN ROUTING
- WORLD-WIDE INTERNET
  - "TELEPHONE" SIZE
  - "INTERNET" TOPOLOGY
- NO MAJOR CHANGE TO ARCH.
  - DEPLOY NOW
  - NOT RESEARCH

HOW DO WE ADMINISTER OSI NSAP ADDRESSES?

- UNIQUENESS
- ADMIN. FEASIBILITY
- ROUTING SCALING
  - => HIERARCHICAL ABBREVIATION
  - EXPECT O(10,000 - 100,000) ~~RDs~~ IN USA ALONE
  - => "FLAT ROUTING" DUBIOUS [OR AT LEAST PAINFUL]

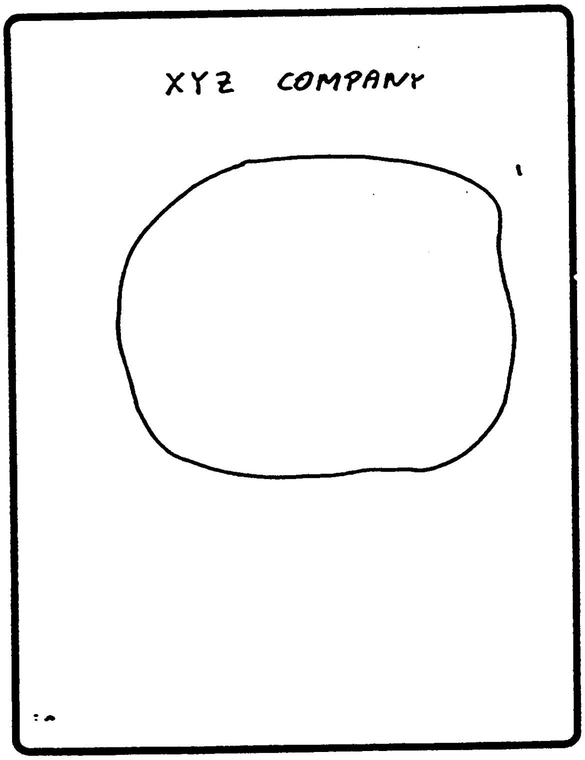
XYZ COMPANY

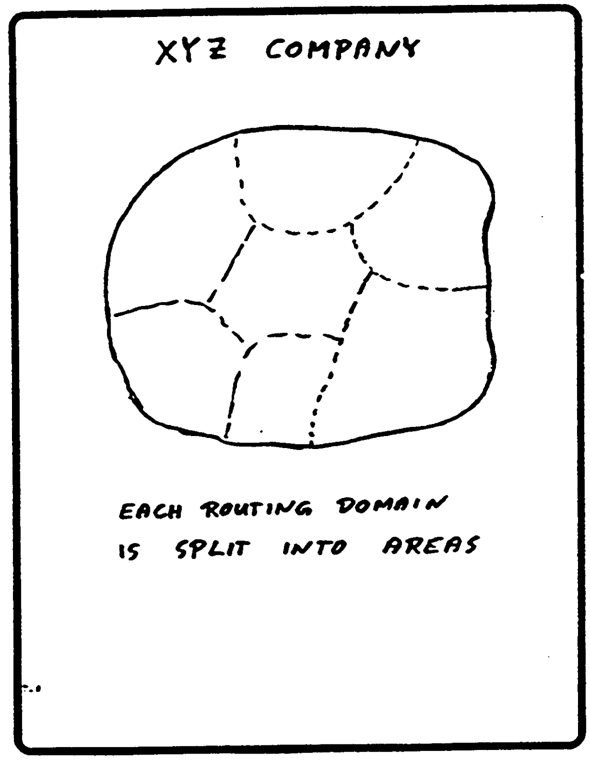
XYZ COMPANY
EACH ROUTING DOMAIN
IS SPLIT INTO AREAS

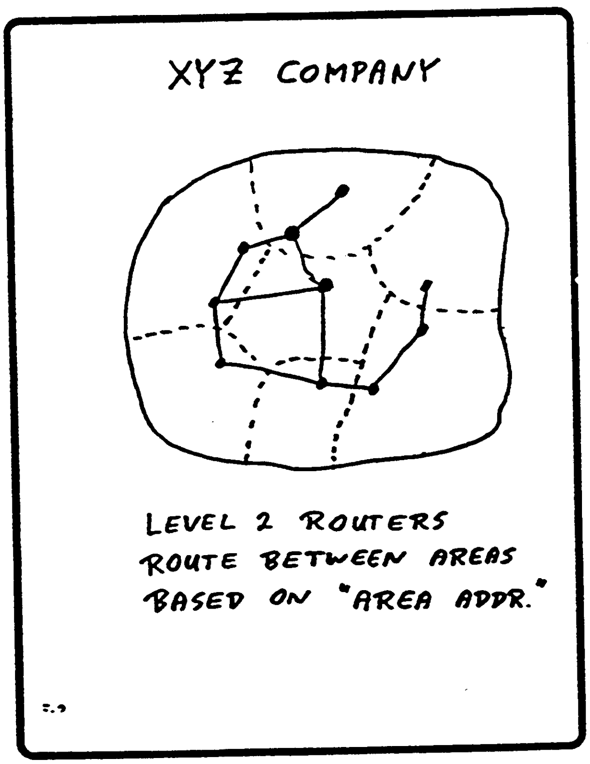
XYZ COMPANY
LEVEL 2 ROUTERS
ROUTE BETWEEN AREAS
BASED ON "AREA ADDR."

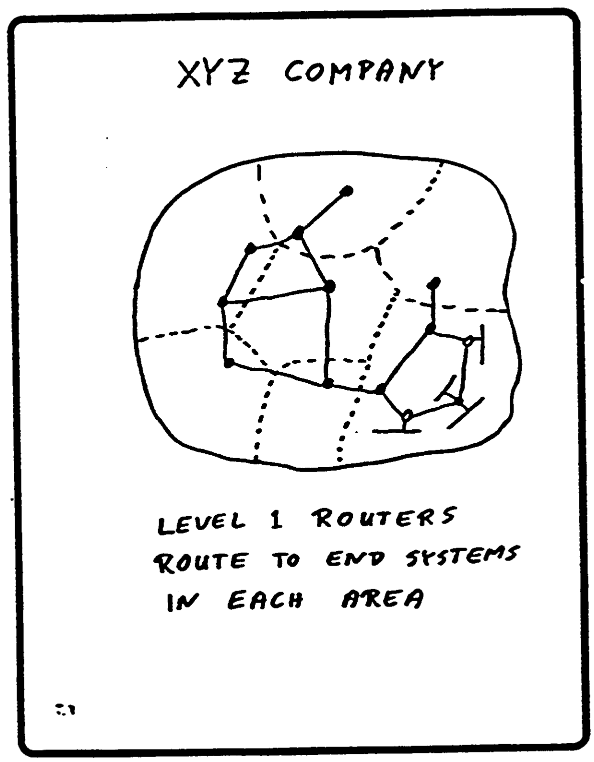
XYZ COMPANY
LEVEL 1 ROUTERS
ROUTE TO END SYSTEMS
IN EACH AREA

XYZ COMPANY

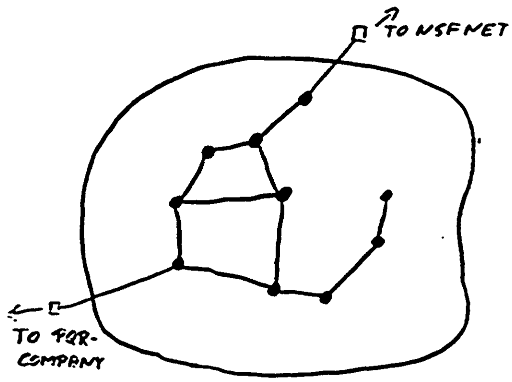

EXTERNAL ROUTING IS BASED ON ADDRESS PREFIX

TRANSIT ROUTING DOMAINS

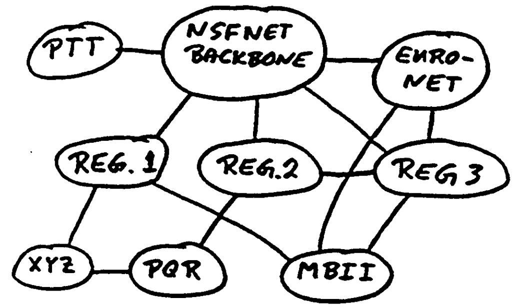

- TRD'S DON'T FOLLOW HIER. MODEL
- STUB DOMAINS MAY ATTACH TO ONE OR MORE (OR ZERO) TRD'S

ADDR => DEFAULT ROUTE

- EG: ORG "O" IS ATTACHED TO REGS "A" + "B"; EITHER
  - (a) GETS ADDR FROM "A" ["USA-A"; "O"]
  - (b) GETS ADDR FROM GSA ["USA-O"]
  - (c) OR, GETS ADDR FROM ISO ["O"]
- IF EVERY TRD IN THE WORLD MAINTAINS ROUTE TO "O"; DOESN'T MATTER

  IF EVERY TRD IN USA MAINTAINS ROUTE TO "O"; (a)+(b) are same

HOW IS-IS USES ADDRESSES

IDP		DSP		
AFI	IDI	HO-DSP	ID	SEL

$\underbrace{\text{AFI IDI HO-DSP}}_{\text{"AREA ADDR."}}$

LEVEL 1: USE "ID" FIELD

LEVEL 2 INTRA-DOMAIN: USE "AREA ADDRESS"

LEVEL 2 INTER-DOMAIN: USE PREFIX

GOSIP/ANSI ADDRESSING

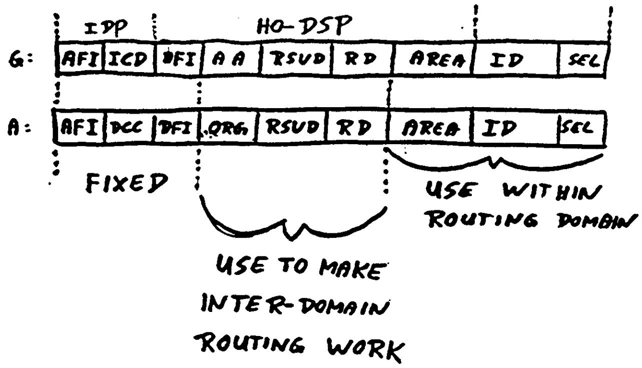

ASSUMPTIONS

- LOTS OF STUB ROUTING DOMAINS (>100,000)
  - CORPORATIONS
  - UNIVERSITIES
  - .....
- MANY ZERO- OR SINGLE-HOMED
- SOME PRIVATE LEASED LINES BETWEEN STUBS
  - NOT AN ISSUE, IGNORE
- A LIMITED NUMBER OF TRD's (10's or 100's)
  - REGIONALS
  - NSFNET BCKBN.
  - AGENCY " S
  - PTT's
- GENERAL INTERCONNECTION OF TRD's

ZERO-HOMED STUBS

- MAY USE PRIVATE LINKS
- ⇒ NEED UNIQUE NSAP PREFIX
- THIS IS EASY
  [GSA + ANSI WILL DO THIS]

SINGLE-HOMED STUBS

- MAY BE A LOT OF THESE
- BIG WIN: ASSIGN THEIR ADDR. ACCORDING TO REGIONAL
  - AA → REGIONAL
  - RD → STUB ROUTING DOMAIN
- HELPS REGIONALS + BACKBONES

SO

⇒ ASSIGN "AA" TO EVERY REGIONAL & EVERY BACKBONE

⇒ HAVE THEM ASSIGN "RD's" TO SINGLE-HOMED DOMAINS

⇒ REGIONALS HAVE TO GET INVOLVED

13

MULTI-HOMED STUBS

- WILL BE SOME
- MIGHT BE LOTS (or LOTS)
- MANY SOLUTIONS
  [EG: COMPANY MBII]
  1. ONE TOP-LEVEL PREFIX
  2. GIVE MBII MULTIPLE PREFIXES (ONE PER TRD). EACH AREA GETS ONE ID BASED ON CLOSEST ATTACHMENT
  3. ONE PREFIX BASED ON ONE OF THE TRD'S
  4. ASSIGN PREFIXES TO SPECIFIC COMBINATIONS OF TRD'S

13½

MULTI-HOMED .....

- CHOICE EFFECTS
  - ROUTING (PATH; BACKUP; ...)
  - RESOURCES (MEMORY, PROTOCOLS, HUMANS, ...)
  - REAL $ $ $ $

⇒ MAY NEED ECONOMIC SOL'N
  - EG: REGIONALS COULD CHARGE FOR "PRIVATE" ADDR'S

⇒ ANY OTHER SOL'N WILL WASTE SOME SORT OF RESOURCE

14

EG: SPECIAL-USE BACKBONE

- CONSORTIUM HAS OWN BACKBONE
  - USED FOR RESTRICTED PURPOSE
  - LIMITED TO KNOWN LIST OF USERS (<1000 SITES)
  - CANNOT BE USED BY "RANDOM GUYS", EVEN TO REACH MEMBERS
    ⇒ DON'T NEED TO ADVERTISE GLOBALLY
- MOST OR ALL MEMBERS MULTI-HOMED
- USE SOL'N #3
  ⇒ EACH MEMBER GETS ADDR. BASED ON OTHER CONNECT.

15

EG: PTT

- SOMEDAY, PTT'S MAY BE UBIQU...
- USE SOL'N #4

EG: TODAYS PHONE SYSTEM (EXCLUDING 800 & 900 & ...)

- IS CLOSEST TO SOL'N #2

6

TRANSITION

- ADDR. BASED ON REGIONAL ⇒ IF REG CHANGES, ADDR. CHANGES
- THIS IS INEVITABLE [IF ORG. CAN "TAKE ADDR ALONG", THEN EVERY TRD HAS TO MAINTAIN ROUTE TO IT].
- WE NEED TO PLAN FOR TRANSITION

7

POLICY QUESTION

- FOR REGIONALS:
  - WILL YOU ACCEPT CUSTOMERS WHO "BRING THEIR OWN ADDR"?
  - WILL YOU CHARGE THEM?
  - WILL YOU ADVERTISE THEM TO BACKBONES ?

  [IF ONLY A FEW, NO PROBLEM]
  [IF MANY, ADVERT. COULD GET BIG]
- FOR BACKBONES
  - DO YOU CARE HOW MANY ADDR. PREFIXES OTHER TRD'S ADVERTISE TO YOU ?

8

CONCLUSION

- ANSI & GSA ARE ASSIGNING NSAP ADDRESSES
- FOR ZERO-HOMED, ALL SET
- REGIONALS & BACKBONES SHOULD BE INVOLVED IN NSAP ADMINISTRATION
  - GET AA FROM GSA or ANSI
  - ASSIGN RD'S TO CLIENTS
- MULTI-HOMED RD'S COMPLEX, HANDLE ON CASE-BY-CASE BASIS
- TRANSITION

9

ALSO: COMMENTS ON NSAP GUIDELINES SOUGHT

COLELLA@OSI3.NCSL.NIST.GOV
CALLON@BIGFUT.ENET.DEC.COM
IETF-OSI-NSAP@OSI3.NCSL.NIST.GOV ←
NOOP@MERIT.EDU

Routing Domains Attached to Argonne

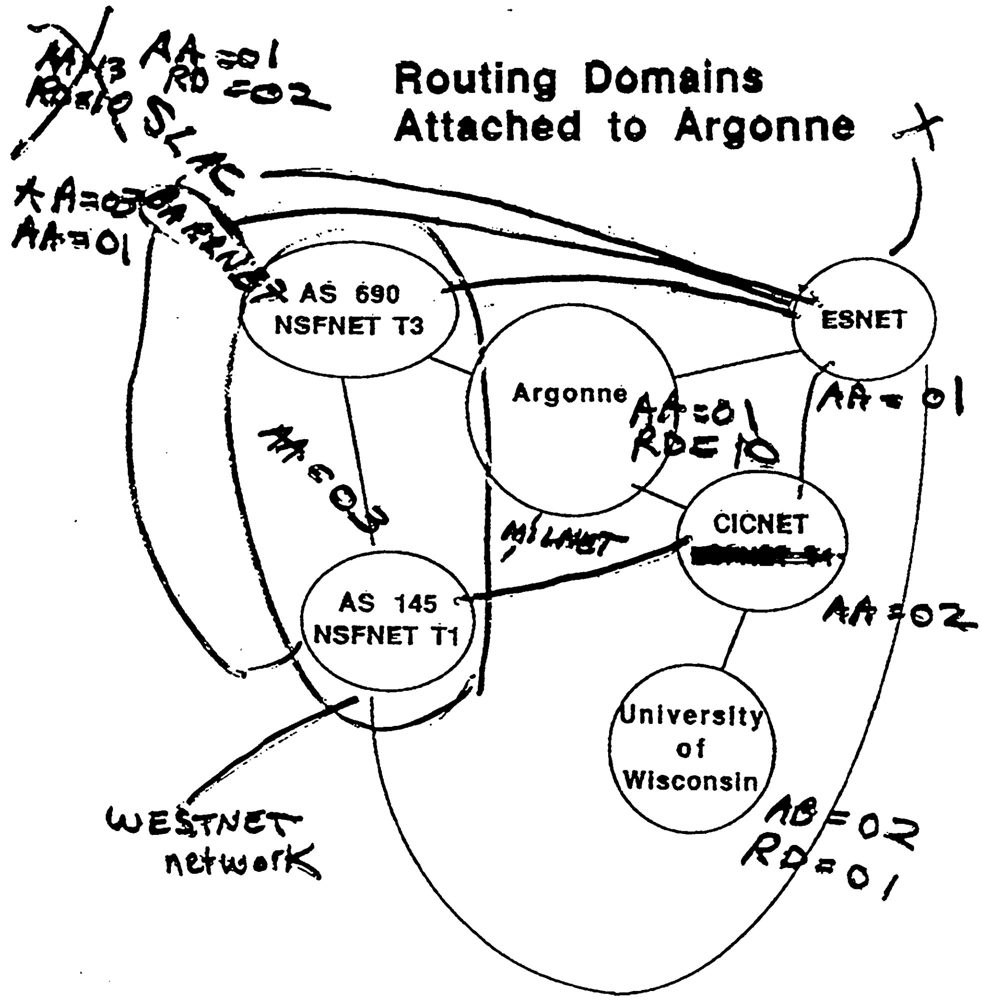

FALLBACK - less optimaL
primary - "BEST" PATH
"Least BAD"

3.4.3 OSI Directory Services (osids)

Charter

Chair(s):
Steve Kille, S.Kille@cs.ucl.ac.uk

Mailing Lists:
General Discussion: ietf-osi-ds@cs.ucl.ac.uk
To Subscribe: ietf-osi-ds-request@cs.ucl.ac.uk

Description of Working Group:

The OSI-DS group works on issues relating to building an OSI Directory Service using X.500 and its deployment on the Internet. Whilst this group is not directly concerned with piloting, the focus is practical, and technical work needed as a pre-requisite to deployment of an open Directory will be considered.

Goals and Milestones:

Mar 1991 Definition of a Technical Framework for Provision of a Directory Infrastructure on the Internet, using X.500. This task may later be broken into subtasks. A series of RFCs will be produced.

Mar 1991 Study the relationship of the OSI Directory to the Domain Name Service.

Ongoing Maintain a Schema for the OSI Directory on the Internet

Ongoing Liaisons should be established as appropriate. In particular: RARE WG3, NIST, CCITT/ISO IEC, North American Directory Forum.

INTERIM MEETING REPORT

Reported by Steve Kille/UCL

OSIDS Minutes

Agenda

- Agenda, Revised
- Minutes of previous meeting
- Liaisons: RARE WG3, NIST, NADF, AARN, ANSI
- Replication
  - Replication Requirements
  - Replication Solutions
  - Network Addresses
  - Presentation Addresses
- APIs for the Pilot
- User Friendly Naming
- Domains and X.500
- Representation of Network Info in X.500
- DSA Naming
- Building Internet Directory/Strategy
- Operational Pilot Status
- Monthly Reports on Pilots
- New Working Groups: Operations, User Support
- Internet Schema
- Naming Guidelines
- Naming for Internet Pilot
- Security
- Directory Assistance Protocol
- Quality of Service
- Date and Venue of next meeting

The meeting was opened by Steve Kille at 9:10am on February 12, 1991. The Agenda was slightly revised and massively reordered. Steve thanked Richard Colella and Peter Whittaker for producing the Minutes. He reported on the status of some of the action items at the last meeting. The formatting of the documents has been improved. The "Infrastructure" document met with some difficulty in forwarding as an RFC. Steve was asked to produce a separate "Strategy" document and to revise the RFC. Steve contacted Al Grimstad to check on a user friendly naming related proposal, and found that this is no longer relevant. There were no corrections to the Minutes.

Liaisons

RARE WG3

Steve reported on this meeting which took place in Brussels in January. They discussed the activities of our IETF-DS group. Their next meeting is April 16-17 in Utrecht, Holland. They meet three times per year. They are very interested in getting more US participation. Future meetings are in July, and also October 31- November 1. Can the IAB find funding for international travel for IETF members? Steve will look into the funding question with appropriate people. European meetings usually have 1-2 representatives from each country. They would also like representation from the FOX project.

NIST

Stuart Cain reported on the Directory SIG meeting in December. They discussed implementation agreements for replication and access control. They would like to see the requirements from our group. NIST is working from the current CDAM. There is already a stable implementors agreement based on the 1988 CCITT recommendation. The new spec is expected by the end of the year. The next meeting will be in March. Steve has replied informally to the NIST liaison to encourage coordination between the two groups and also to share our documents on replication requirements and solution. The sense of this was agreed to by the group, and it will be used to generate a formal liaison response. The NIST group is concerned with "freezing" their agreements based on a DIS version of the standard, and will be working to avoid that kind of discrepancy.

North American Directory Forum

Marshall Rose reported that the last meeting was in October, before the last IETF-DS meeting. The next meeting is in March, after this meeting.

Australian Academic Research Network

Steve received a liaison statement from George Michaelson. Australia is working on X.500 naming and addressing standards. They will send people to the IETF some time this year. They have not been able to participate in this group due to lack of funds.

ANSI US Directory Ad Hoc Group

Roy Van Dorn (HP) reported that this group met last week. They are bringing ballot comments to ISO. Subordinate references will be replicated, according to the latest draft standard. Replicating cross-references will not occur. Hoyt Kesterson is the ISO Rapporteur. Skip Sloan will be the head of the US delegation. Steve will send them the replication documents from our group. There will be one more US meeting in March for ballot comments. The liaison of the group's documents to ISO will be done through ANSI by Paul Koski. Access control and replication are US priorities. Some of the schema document will get into the 1992 standard. The definitions of attributes will be more like 1988. The four types of object classes will continue. Subtrees and partial entries within subtrees can be replicated. A completeness flag is included in replication. Searches on attributes that don't exist will be referred for further lookup. The unit of replication is an entry, not an attribute within an entry.

- Replication
  - Replication Requirements

    It was agreed that this Internet Draft (Replication Requirements to Provide an Internet Directory Using X.500) be progressed to an RFC.

  - Replication Solutions

    There was substantial discussion of this paper. Marshall and Steve revised the text during the meeting and redistributed the document. Marshall suggested that the title be changed to include the changes to Distributed Operations as well as replication. This suggestion was agreed to by all. A number of changes were suggested to make the document more clear. There was a suggestion to include a figure describing knowledge replication. None of the proposed changes require discussion at a further meeting, and Steve agreed to send a revised document out to the list on Monday (February 18). The group will respond within one week with any comments. After that the Internet Draft (Replication to Provide an Internet Directory Using X.500: A Proposed Solution. However the title may be changed.) will be progressed to an RFC.

  - Network Addresses

    There were a few comments from the IAB regarding the Telex kludge. It was agreed that this Internet Draft (An Interim Approach to Use of Network Addresses) be progressed to an RFC.

  - Presentation Addresses

    It was agreed that this Internet Draft (A String Encoding of Presentation Addresses) be progressed to an RFC.

- APIs for the Pilot

  Ruth Lang said that this was an important area and would like to see suggestions for APIs (application programming interface). The only comment received so far on the list was from Peter Whittaker (BNR) about object management support in XOPEN. There was a discussion of the XDS agreements. Peter Mierswa said that DEC participated in XDS. The user-friendly and object-oriented aspects of XDS will cause applications to be large. It is difficult to extend the XDS object set. There are other technical drawbacks, but it was agreed to by a number of parties. DEC will support the XDS API but also a more functional layer. Quipu does not support XDS. XDS and object management documentation is available from Omnicom. It was felt that APIs did not fit into our group's charter. We may want to make recommendations but then move on to the technical infrastructure. This group is also not to manage projects or pilots.

- User Friendly Naming

  Peter Mierswa tried to find a common syntax set with the OSF DCE naming (based on unix filesystem syntax) and the proposed X.400 annex for business card OR address format (uses semicolons and slashes, which evolved out of the RFC 987 work). However there was no such syntax in common and Peter gave up. The algorithm in this document is useful based on experience, though there may be scope for experimentation. It was noted that name space organization affects efficiency of searches. For example Cambridge University uses many levels of OU. It is recommended in the Naming Guidelines document (see section 18) that pilots be laid out so that this user friendly naming scheme works reasonably. It was agreed that this Internet Draft (Using the OSI Directory to Achieve User Friendly Naming) be progressed to an RFC.

- Domains and X.500

  UCL has done some work in implementing this scheme. There is a tool to do a white pages lookup based on a domain address. This is an experimental service. The general appropriateness of representing domain name system information in the Directory was discussed. This is viewed as controversial. The X.500 version of DNS may have be usable for other functions than those currently offered by the DNS, such as browsing. Mailbox records are included in the DNS, but are not widely used. Peter Mierswa said that it would not matter if this was not submitted as an RFC. Steve disagreed with that and would like to progress the work. Tim Howes suggested that we submit this with a disclaimer that it is experimental. Steve would like the IAB to discuss these issues. Jose Garcia-Luna felt that security should be discussed in this paper. It was eventually agreed that this Internet Draft (Domains and X.500) should be progressed as an RFC.

- Representation of Network Information in X.500

  Mark Knopper and Chris Weider gave a presentation on some work in progress at Merit, which will become part of the DARPA/NSF sponsored Field Operational X.500 (FOX) project. They have entered the network contacts part of the whois data into the @o=Internet part of the White Pages DIT. New object classes have been defined. Bill Nowicki noted that putting all of the IP network numbers into a single location in the DIT will not scale well. It was suggested that the network number entries be located within the owning organizations. This would obviously require much more participation in the X.500 projects. For now the net numbers can be entered in a separate tree under o=Internet and eventually these entries will just be pointers to the master network entries. Steve proposes another solution to this in the Domains and X.500 paper. It is scalable, but also requires more work to bootstrap. There will be further cooperation with SRI, ISI and PSI to allow the rest of the NIC's data to be entered into X.500. There were a number of useful suggestions on how the network information could be stored in the DIT. It was recommended that Merit produce an internet draft to document this effort, both work in progress as well as long term design. Chris agreed to do this by March 7. He will take the scalability issues into account.

- DSA Naming

  The current South American wildlife names don't seem to be descriptive enough! The solutions outlined in this paper solve some operational problems with quipu-based pilots. Peter pointed out that the section on multinational organizations does not solve the problem. There were several suggestions for modifications, and discussion of this will be necessary at the next Working Group meeting. It was felt that after that, this Internet Draft (DSA Naming) can be progressed to an RFC.

- Building Internet Directory/Strategy

  The infrastructure Internet Draft was held up in protracted discussion regarding how to submit RFCs. Steve wrote a new strategy document. It was agreed that APIs should be mentioned in this document. The "strategy" was removed from the I.D. and so that was renamed to a very long name beginning with "Overall Plan". The strategy document was agreed to in principle but will not be forwarded at this time. The Overall Plan Internet Draft was agreed to be progressed to an RFC again.

- Operational Pilot Status

  - PSI Pilot

    Marshall reported that there are about 70 organizations on the US pilot. Growth has been linear since the pilot began. ISODE 6.8 interim release is due out by the end of the month. It is a very stable and higher performance version. It will have Tim Howes' mods to quipu, and also the Directory Assistance Protocol (which allows splitting the DUA between two different hosts). FRED is faster now. There is a Macintosh DUA offered by PSI as shareware. A source license is available similar to the Nysernet SNMP license. The PSI pilot only allows DSAs to be connected via IP (and now CLNP). The quality of X.25 in the US "provides pneumatic inward pork-pressure via narrow flexible tubing."

  - COSINE Pilots

    Steve reported that 19 out of 20 countries in COSINE are running X.500 pilots. The COSINE P2.1 pilot has been renamed as PARADISE, and has officially started. Its manager is David Goodman. ULCC has an operational facility to replace Giant Tortoise. Their plan is to support international pilots until the end of 1992. France has a research pilot based on quipu and also a commercial pilot based on Pizarro. Xtel and the Dutch PTT are involved in PARADISE.

- Monthly Reports on Pilots

  It is felt that the operational pilots should distribute status reports on a monthly basis. The FOX project is interested in coordinating the US report. Ruth Lang contacted Jon Postel at ISI about this and Jon volunteered ISI to produce the reports. Some FOX mailing lists will be set up to help coordinate the US report. David Goodman, the PARADISE manager, will integrate this into the international report. FOX and PARADISE will agree on timescales for ensuring that this comes out each month.

Reports will be timely, with noncontributors marked as "no report for XXX". This international report will be sent out as a part of the Internet Monthly Report and to a separate list for those not interested in other aspects of the IMR. The reports should be on "The State of the DIT". Organizations should be queried for their activities. Marshall gets regular statistics reports from the US DSAs. The Canadian pilot is operated by the University of Toronto.

- New Working Groups

  - X.500 User Support Working Group

    Chris Weider volunteered to Chair a new Working Group. Steve will talk to the IETF area coordinators and suggest that the new group be jointly in the OSI and User Services areas. Several of the group participants were interested in joining the new group. The first meeting will be at the next IETF. Chris distributed a draft charter and several comments were made. Chris will talk to Joyce Reynolds and Dana Sitzler, to see whether it would be reasonable to model the group after the NISI Working Group. Perhaps the new group should be called DISI (pronounced "dizzy"). The group would provide a documentation package for sites, as well as a center of expertise for X.500 issues.

  - X.500 Operations Working Group

    There was some interest in forming such a group but it was felt that this should wait until the activities of the main IETF-DS group come to an end, or at least go into "maintenance mode". It was viewed that the group will only last for one more meeting with the same high level of activity. After that the operations group will be formed. Marshall Rose and Chris Weider were involved in discussing the charter of the new group.

- Internet Schema

  Marshall suggested that the name of the Internet Draft (COSINE and Internet Naming Architecture) be changed from "naming architecture" to "schema". This was accepted. There were comments on this document at the RARE WG3 meeting. The textEncodedORAddress attribute was deprecated by OSI purists, but some members felt it was useful in the pilots. This Internet Draft was agreed to be progressed to an RFC.

- Naming Guidelines

  Steve introduced this Internet Draft and explained that it sets out some guidelines for how to lay out a pilot DIT. It is a follow-on to annex B of X.521. Marshall mentioned that the T.61 character sets for international symbols once were a problem but work now in quipu. Peter mentioned that this is not a solution for multinational organizations. It is viewed that this is a difficult problem, and that the acceptable solutions should be documented. There needs to be a definition of "multinational organization". HP would like to see a single "mount point". There was a discussion of organization naming strategy. Marshall suggested that the names be fully descriptive to avoid later, possibly legal, conflicts. The naming authorities must enforce unique names within the DMD. Long names were recommended. Marshall mentioned that a small DIT depth makes browsing less effective. It is not useful to define conformance rules for a guidelines document. Conformance is useful for a given national pilot. Steve and Paul Barker will edit the document and distribute to the group. At the next meeting it will be proposed that the Internet Draft (Naming Guidelines for Directory Pilots) be progressed to an RFC.

- Naming for Internet Pilot

  Marshall gave a presentation of a paper he and Einar Stefferud had written to be presented at the NADF, US-CCITT-Study Group D, and ANSI as well as to this group. The problem is that there are no OSI numbering authorities in the US, but they are needed for pilots to advance to a production stage. ANSI has accepted over 500 applications for OIDs under 1.2.840, but due to legal problems have not assigned any. Numbers are not a problem for ANSI but names are. The only legal method would be to assign the name and then publish the fact in the Federal Register with the reserve to revoke on a 6-month challenge procedure basis. GSA has been assigning NSAPs under AFI/IDI=47/0005, only for federal agencies. IANA has assigned several hundred OIDs under 1.3.6.1.4.1 for internet network management use. US-CCITT-SG-D is trying to make a national decision on naming, but only for an X.400 ADMD/PRMD registry and not for X.500. Possible naming universes are geographical, political or community. Civil authorities are the best choice as it gives a familiar and undisputed structure. However collisions in RDNs must be avoided. The proposal suggests using the numeric code assigned by ANSI for the RDN itself. This was heavily disputed, but as Marshall noted it would be legally defensible. The consensus was that we should fix ANSI rather than using numeric RDNs. Marshall and Stef believe that their presenting this proposal to the four groups will force a national decision. The proposal went on to recommend use of numeric codes for states and populated places. Naming of OSI entities was included, and there was a suggestion that non-OSI entities should get names too (e.g., SNA, TCP/IP applications). Steve suggested that this be made into an Internet Draft but not a standard. Marshall will make the changes suggested by the group before the NADF presentation in March. He will "lean heavily" on ANSI to begin assigning names. Beth Summerville is ANSI's registrar for the naming authority function.

- Security

  Peter Yee's paper was revised since the last meeting. There were not many changes due to lack of comments at Boulder. Marshall said that it will be necessary to consult with the IETF Security Working Group before progressing this document. Peter will contact Steve Crocker to get help on proper security terms and concepts. Marshall suggested splitting the discussion in the paper between authentication (simple now, strong later), and authorization (access control lists). Paul suggested including an ACL to control access for searching. Steve suggested that this should become an Internet Draft with title Security Requirements for X.500 in the Internet. There should be a companion document for Security Solutions, and this should reference the 1992 CCITT document. A problem at MIT is that they want to limit searching their organizations to return data only if less than n entries. HP wants to disallow searching their organization entirely. Peter will revise the document and send it out to the list by March 1.

- Directory Assistance Protocol

  Marshall wrote an RFC describing a protocol used by PSI's Macintosh DUA client. It documents existing practice and is not a standard. The server is part of ISODE. He characterized the protocol as "horrid". Tim Howes has also been working on a Macintosh DUA with a different protocol. Tim will write an RFC for his DAP pretty soon.

- Quality of Service

  Steve submitted an informal writeup to suggest that QOS attributes be added to the schema to represent the advertised quality of DSA services in the pilots. This was thought to be a good idea and there were no objections to including this in the Schema document.

- Notable Actions, Dispositions and Promises

  - RFC Progression

    The following documents were recommended to be progressed to RFC status:

    Replication Requirements to Provide an Internet Directory Using X.500 (section 6a)

    Replication Solution and Distributed Operations (section 6b)

    An Interim Approach to Use of Network Addresses (section 6c)

    A String Encoding of Presentation Addresses (section 6d)

    Using the OSI Directory to Achieve User-Friendly Naming (section 8)

    Domains and X.500 (section 9)

    Overall Plan (section 12)

    Internet Schema (section 16, and including QOS item in section 21)

    Naming Guidelines for Directory Pilots (section 17)

  - Action Items

    Strategy document will be revised by Steve (sections 4, 12). The issue of travel funding will be investigated by Steve (section 5a). A formal response to NIST will be drafted by Steve (section 5b). The replication documents will be sent to ISO via ANSI and Paul Koski by Steve (section 5c). Jon Postel, for the FOX project, will set up a mailing list, and produce monthly reports coordinated with PARADISE and the Internet Monthly Reports (sections 10 and 14). Chris Weider will start up the new Directory Information Services Infrastructure Working Group (section 15a). Chris and Mark will write an RFC on representing network infrastructure information by March 7 (section 10). Marshall Rose will lean heavily on ANSI to assign organization ids and names (section 18). The security document will be revised by March 1 by Peter Yee (section 19).

- Date and Venue of Next Meeting

  There will be no OSI-DS meeting at the March IETF. The next meeting will be after that, to be decided on the list. A possibility is a video conference, or alternatively a face to face meeting either in Ann Arbor or on the east coast in May or June. The choice depends on online discussion of the working drafts. Given some comments, it might be appropriate to wait until July. Steve will poll the group after the next round of editing.

- Thanking the Host

  Ruth Lang and SRI International were thanked for their excellent services including a lunch.

Attendees

Stuart Cain	scain@hpindeg.cup.hp.com
Cyrus Chow	cchow@orion.arc.nasa.gov
J.J. Garcia-Luna	garcia@sri.com
Arlene Getchell	getchell@nersc.gov
Barry Holroyd	berries@eng.sun.com
Tim Howes	tim@d.cc.umich.edu
Steve Kille	S.Kille@cs.ucl.ac.uk
Mark Knopper	mak@merit.edu
Paul Koski	koski@hpindeg.cup.hp.com
Ruth Lang	rlang@nisc.sri.com
Peter Mierswa	mierswa@smaug.enet.dec.com
Bill Nowicki	nowicki@sun.com
Alex Pepple	alexp@bnr.ca
Marshall Rose	mrose@psi.com
Chris Weider	clw@merit.edu
Russ Wright	wright@lbl.gov
Peter Yee	yee@ames.arc.nasa.gov

3.4.4 OSI General (osigen)

Charter

Chair(s):

Robert Hagens, hagens@cs.wisc.edu
Ross Callon, callon@bigfut.enet.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.

3.4.5 OSI X.400 (osix400)

Charter

Chair(s):

Rob Hagens, hagens@cs.wisc.edu

Mailing Lists:

General Discussion: ietf-osi-x400@cs.wisc.edu
To Subscribe: ietf-osi-x400-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.

3.4.6 Office Document Architecture (oda)

Charter

Chair(s):
Peter Kirstein, kirstein@cs.ucl.ac.uk

Mailing Lists:
General Discussion: ietf-osi-oda@cs.ucl.ac.uk
To Subscribe: ietf-osi-oda-request@cs.ucl.ac.uk

Description of Working Group:

The ODA Working Group will develop guidelines for the use of the Office Document Architecture for the exchange of Compound documents including formattable text, bit-map graphics and geometric graphics according to the ODA Standard. It will consider also Intercept Standards for other document content types it considers vital - e.g., Spreadsheets. The Working Group will define how to use both SMTP and X.400 for interchange of ODA documents. It will maintain close liason with the SMTP and X.400 Working Groups.

This Working Group will review the availability of ODA implementations, in order to mount a Pilot Testbed for processable compound document interchange. Finally, it will set up and evaluate such a testbed.

Goals and Milestones:

Done	Inaugural meeting.
Jul 1991	Produce a paper stating what ODA standards or profiles still need completing.
Jul 1991	Produce paper on how both SMTP and X.400 message systems should be supported.
Jul 1991	Produce paper on what pilot implementations can be provided.
Jul 1991	Produce paper on what scale and type of Pilot Testbed should be organised.
Dec 1991	Provide first feedback on the ODA Pilot.
Ongoing	Coordinate ODA Pilot.
Ongoing	Review and propose additional enhancements of ODA.

CURRENT MEETING REPORT

Reported by Peter Kirstein/UCL

ODA Minutes

Agenda

- Introduction of Participants.
- Discussion of Charter.
- Review of Documentation available.
- Consideration of current status of standardisation.
- Review of facilities needed for IETF-ODA Pilots.
- Review of possible products.
- Discussion of interaction with message systems.
- Discussion of interaction with other Working Groups.
- Review of possible programme and timetable.
- Proposed further actions.
- Methods of Working.
- Arrangements for future Meetings.

The attendees outlined their interests in the Working Group. Most were interested to use facilities provided to them; few were interested in developing facilities themselves. There was interest in the functionality of ODA, therefore a tutorial by Fred Held was organised as an evening session; it was attended by about 25 people.

The group agreed that they would like to use existing software - but needed to know what was available.

The Chair outlined the capabilities of ODA; it would enable the interchange of documents with various text capabilities (including Fonts), geometric graphics and bit-map graphics. It would allow, therefore, interchange of processable documents between different word processors. The bit-map graphics supported both Group 3 and Group 4 facsimile formats - potentially of interest to the NETFAX Group. The standard is very general. To ensure the capability of document interchange, it is essential to define also a Document Application Profile (DAP), to which any product must conform. A particular DAP has been developed in Europe under the PODA project, and a number of products exist to this DAP (Q112,[1]). As background for the discussions of the ODA Working Group, some papers have been made available in an electronic form. These are listed below, and may be requested from the UCL-cs info-server: "info-server@cs.ucl.ac.uk".

The documents are accessed by standard message systems, giving a message body of the form:

request:ietf-osi-oda topic:xxxx

where xxxx is the name of the document required.

Currently a number of documents are available in the info-server - all in text form. Many could be available in ODA/ODIF format if this was required. The documents are listed in a document called INDEX in the collection.

The Chair stated that software will be available to allow documents preparation and storage, and also document interchange to the DAP. He had identified three products which would support ODA from the beginning at the 3rd quarter of 1991: the SLATE editor from BBN (with UCL additions), a product from Xerox, and various DEC products for CDA. A version of WORD from Honeywell-Bull, and of WordPerfect from ICL would probably exist, and other products could be available by the summer. It was proposed, and agreed, that the group will try to get started as soon as possible on a pilot activity. The members of the group would want to experiment with the facilities themselves; if they were satisfactory, they could try to get other user groups interested.

For a User Pilot, it was necessary to have not only an editor which could produce an ODA stream (ODIF), but also combine it with a mail system. The ODIF stream contained arbitrary 8 bit binary; therefore it could not be sent by RFC 822 mail without modification. Luckily the SMTPEXT group was proposing both a short-term and longer term recommendation for the extension of that system to support binary data. Another mail system (X.400) was the brief of the OSIX.400 Working Group; that system also supported binary data. It was agreed that the present Working Group make known its needs to, and use the mail systems defined by, the other two Working Groups. We need not consider mail further inside the present Working Group - except to make recommendations based on the actions received from the other groups. Some of the products of interest with the ODA capability (WORD, WordPerfect) existed currently only for PCs. The Working Group participants felt that they were already making adequate ad-hoc arrangements to incorporate documents from PCs into mail systems, and did not need - or want - the Working Group to address the mechanisms needed.

In accordance with the Charter, the Chair promised to provide further details of product availability before the end of April. By that time, the interim recommendation of the SMTPEXT Working Group should be available. The aim was still that sufficient information should be available by that time, that an initial set of trials by participants should be possible between the lst and 2nd quarters of 1991, and that a detailed plan for a PILOT should be ready for the next IETF meeting in Atlanta.

It was not thought necessary to have a further meeting prior to the next IETF, but a meeting during that week was planned.

A set of documents relating to ODA had been put in an archive - further documents will be added to this database as they become available.

Reference

1. EWOS: ODA Document Application Profile Q112 - Processable and formatted documents - Extended mixed mode, PrENV 41 510, Paris, 1988.

Attendees

Richard Bowles	bowles@stsci.edu
Ross Callon	callon@bigfut.enet.dec.com
David Crocker	david_crocker@palo-alto.pa.dec.com
Shari Galitzer	shari@gateway.mitre.org
Russ Hobby	rdhobby@ucdavis.edu
Darren Kinley	kinley@crim.ca
Peter Kirstein	kirstein@cs.ucl.ac.uk
Jim Knowles	jknowles@trident.arc.nasa.gov
Vincent Lau	vlau@sun.com
David Miller	dtm@ulana.mitre.org
Robert Morgan	morgan@jessica.stanford.edu
Robert Reschly	reschly@brl.mil
George Sanderson	sanderson@mdc.com
Mark Sherman	mss+@andrew.cmu.edu
Gregory Vaudreuil	gvaudre@nri.reston.va.us
John Veizades	veizades@apple.com
Wengyik Yeong	yeongw@psi.com

3.4.7 X.400 Operations (x400ops)

Charter

Chair(s):

Alf Hansen, Alf.Hansen@pilot.cs.wisc.edu

Mailing Lists:

General Discussion: ietf-osi-x400ops@pilot.cs.wisc.edu
To Subscribe: ietf-osi-x400ops-request@pilot.cs.wisc.edu

Description of Working Group:

X.400 management domains are being deployed today on the Internet. There is a need for coordination of the various efforts to insure that they can interoperate and collectively provide an Internet-wide X.400 message transfer service connected to the existing Internet mail service. The overall goal of this group is to insure interoperability between Internet X.400 management domains and to the existing Internet mail service. The specific task of this group is to produce a document that specifies the requirements and conventions of operational Internet PRMDs.

Goals and Milestones:

Done	Initial meeting, produce internal outline.
Mar 1991	Working draft, circulate to interested people.
Jul 1991	Internet Draft available.
Dec 1991	Document ready for publication.

CURRENT MEETING REPORT

Reported by Kevin Jordan/CDC

X400OPS Minutes

Review of Agenda

The Agenda was approved without change, although, some minor adjustments were made as the meeting progressed.

Review of Charter

It was made clear that the focus of the Working Group is the operation of X.400 mail on the Internet.

Rob Hagens presented a one page draft document describing the strategy for deployment of X.400 in the Internet. The goals described in the document were reviewed and discussed. The goals drafted by Rob were:

- The X.400 service will not, in the near future, completely replace existing Internet mail service.
  - It was pointed out that this is an assumption, not a goal. It was suggested that a useful goal would be to work with the SMTPEXT Working Group in order to facilitate gatewaying between SMTPEXT and X.400.
  - People who had attended the SMTPEXT meetings on Monday, 3/11, reported briefly on what was discussed there. It seems that the SMTPEXT Working Group has just begun defining requirements, so judging from previous experience, it will probably be at least two years before SMTP V2 is widely implemented and operational.
- The X.400 service in the Internet shall be fully connected to the existing Internet Mail service via gateways.
  - It was recommended that this goal be revised so that it includes a clause about the need for X.400 gateways to be highly interoperable with the existing Internet mail services.
- The X.400 service in the Internet will be connected to international R&D X.400 service initiatives.
  - UW-Madison has already established a direct X.400 link to Norway, Finland, Canada, UK, France, Switzerland and Spain. The Norwegian connection has agreed (at least for now) to act as a relay between XNREN and the other participants of the COSINE X.400 project in Europe, not directly connected to XNREN.

- The X.400 service in the Internet will be connected to major ADMD providers in the U.S., provided that a suitable arrangement can be made.
  - There was general consensus that this is a very important goal. However, it is not yet clear how this goal will be attained due to the fact that the ADMD providers are commercial organizations who normally account and charge money for their services.
  - On the second day of meetings, Vint Cerf indicated that CNRI is already pursuing this goal. CNRI is willing to provide the physical plant necessary to provide a connection to an ADMD provider, but human resource limitations may delay implementation. Rob Hagens indicated that UW-Madison could help.
- Although the 1984 protocols may be used on an experimental basis, the primary deployment of X.400 should be based upon the 1988 version of X.400.
  - It was recommended that this goal should be rewritten in terms of driving toward general deployment of 1988 X.400 (or perhaps 1992 X.400), but that it is also necessary to provide backward interworking with 1984 X.400. Conversion from 1984 to 1988 to 1992 and beyond will not occur all at once. The transitions will probably be gradual, so backward interworking is desirable.
- With respect to management domains, the Internet will be organized as a collection of Private Management Domains.

Finally, the Technology section of the draft document contained the following statement:

> The X.400 service in the Internet will conform to the US GOSIP profiles.

It was recommended that this statement be qualified because, for example, GOSIP requires OSI lower layers, but the Internet X.400 service will be based primarily upon TCP/IP (RFC1006) initially.

Relationship to other technical groups

Some members of the X.400 Operations Working Group are also members of other technical groups. It was suggested that this informal cross participation would be used for communications between the X.400 Operations group and other groups. The groups mentioned were: IETF-DS, IETF-ODA, RARE-WG1, R&D MHS Managers, NIST Workshops.

Round table presentation of current X.400 service status

Each of the Working Group participants discussed how X.400 is being used (or is planned to be used) within his/her organization. Most sites are planning to use X.400, but are not using it actively yet. Notable exceptions are UW-Madison and CDC; these organizations are actively using X.400 now.

Overall organization of the X.400 service in the Internet

- Technical requirements

  Two types of MTA's were defined:

  - MTA's supporting RFC1006, informally called Internet MTA's
  - MTA's supporting TP0/X.25, informally called PDN MTA's

  It was generally agreed that organizations wishing to participate in the Internet X.400 project should support Internet MTA's, meaning that participating organizations should provide an MTA which supports RFC1006.

  However, the Working Group does not want to preclude participation by organizations which are connected only to X.25-based PDN's. Such an organization will need to make a bilateral agreement with an organization which supports both RFC1006 and TP0/X.25, and arrange for that organization to relay mail between the X.25-based connection and the RFC1006-based Internet connection, or each PRMD should implement mechanisms to insure end-user connectivity on top of both stacks.

  We should also be prepared to serve MTA's connected to the TP4/CLNP infrastructure.

  It was noted that these technical requirements are essentially the inverse of the connection requirements established by COSINE for its members. COSINE requires all participating organizations to support TP0/X.25 connections to their respective country's PDN. RFC1006 is not defined as mandatory by COSINE. This implies that interconnection of COSINE and the Internet X.400 project will:

  - Require a relay in the U.S. to support both X.25 and RFC1006, or

  - Require a relay in Europe to support both X.25 and RFC1006. This, in fact, is the current state of affairs, or

  - Combinations of a. and b. above.

  It was generally agreed that GOSIP should serve as a reference document for X.400 upper layer technical requirements, where "upper layers" is defined to be the OSI Session layer and the layers above it.

  The term "Internet WEP" was introduced to identify a special MTA acting as a Well-Known-Entry-Point for an Internet PRMD. UW-Madison will distribute a draft definition of an Internet WEP to the list for review.

- Internal organization of PRMD's

  It was agreed that naming authority should be hierarchically organized. Specifically, the names of organizations should be coordinated with the PRMD's in which the organizations are created. Similarly, the names of organizational units should be coordinated with the organizations in which the organizational units are created (but not necessarily with the PRMD administrators).

UW-Madison will maintain a list of Internet PRMD's.

UW-Madison will maintain FTP-able documents which describe participating organizations and information about MTA's (e.g., MTA connection information). ONLY operational organizations and MTA's will be described in these documents.

It was agreed that an important characteristic to describe about an MTA is its ability to operate over both RFC1006 and TP0/X.25. Publishing this characteristic will make it easy for prospective participants supporting only TP0/X.25 to locate existing participants who might be willing to act as Internet relays.

UW-Madison will distribute a draft definition of an MTA document format to the list for review.

Specification of RFC822 addresses in the X.400 world

It was agreed that RFC822 addresses should be expressed using X.400 domain defined attributes. Furthermore, a special PRMD named "Internet" will be defined to facilitate the specification of RFC822 addresses. For example, an X.400 user will address an RFC822 recipient by constructing an X.400 address such as:

/c=us/admd= /prmd=Internet/dd.RFC-822=user(a)some.place.edu/

Participating MTA's will be configured to recognize "/c=us/admd= /prmd=Internet/" as a special case. The presence of this address will cause a message to be routed to a regional RFC987 gateway. In effect, this special PRMD identifies a community of gateways to RFC822 recipients. This strategy is user friendly in that all users everywhere need only remember this one gateway address, and it is efficient in that it avoids having to establish a single, common gateway which would tend to become a bottleneck and single point of failure.

Specification of X.400 addresses in the RFC822 world

After considerable discussion, it was agreed that RFC822 users should be able to address X.400 recipients in RFC822/Internet terms. This implies the necessity of maintaining and distributing address mapping tables to all participating RFC987 gateways, at least in the short-term. Other mapping strategies were discussed (loudly and enthusiastically), but it was shown that these alternate strategies would sometimes cause messages (or replies to messages) to pass through more than one gateway. Since this behavior would probably cause information to be lost in translation, it was quickly agreed that the alternate strategies were inferior to the good old table driven approach.

Nevertheless, it was also pointed out that some X.400 addresses do not map cleanly to RFC822 addresses, even when the table driven mapping strategy is used. For example, X.400 personal names which contain generation qualifiers, personal names which contain initials but no given name, and initials which contain periods cannot be mapped to RFC822 symmetrically such that a reverse mapping is possible. Similarly, X.400 addresses which contain X.121 address elements (sometimes used for expressing fax telephone numbers), unique UA identifiers, or physical addressing attributes cannot be mapped nicely. Consequently, it will be necessary for RFC987 gateways to generate RFC987 address syntax occasionally.

It was recommended that our RFC should contain guidelines for the creation of X.400 personal names. In following these guidelines, users will avoid creating personal names which can not be mapped nicely between X.400 and RFC822.

It was generally agreed that long-term reliance upon static mapping tables is unacceptable. Therefore, it was agreed that the X.400 Operations Working Group should devise a strategy for using X.500 directory services instead.

Another option could be to use the DNS system for this purpose, if the X.500 infrastructure appears to be too premature.

Future issues

The following list of issues were agreed to be important for the future service, and the group should follow these issues closely:

- X.400/84 <--> 88 interworking.
- Use of DNS for RFC 987 address mapping management
- Use of an X.500 infrastructure for routing, table management and user catalog purposes.
- Body types other than text.

Presentation of outline for RFC

Rob Hagens proposed an outline for the RFC to be produced by the Working Group. Participants made comments and suggested additions.

UW-Madison will write a first draft and distribute it to the list for review.

Future meetings

A tentative meeting has been scheduled for May 30 and 31. This meeting will be held in Madison, Wisconsin or San Jose, California. The purpose of the meeting will be to resolve comments against the draft RFC, in case there are comments which can not be resolved via email.

The next general IETF meeting is scheduled for July 29 - August 2 in Atlanta, Georgia.

The X.400 Operations Working Group will definitely meet at that time.

Action items

Rob Hagens	Update and distribute the X.400 Internet Service Strategy document.
	Update and distribute outline for the RFC.
Alf Hansen	Write and distribute a definition of "Internet WEP".
Kevin Jordan	Distribute Minutes from St. Louis meeting. Distribute the X.500 schemas used by CDC to record information about X.400 routes, MTA's, and address mappings. Include a description of how these schemas are used by CDC's X.400 products.
	Distribute a description of CDC's extensions to RFC987 in support of multipart/multimedia X.400 messages.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Vinton Cerf	vcerf@NRI.Reston.VA.US
Cyrus Chow	cchow@orion.arc.nasa.gov
Alan Clegg	abc@concert.net
John Dudeck	jdudeck@polyslo.calpoly.edu
Ned Freed	net@ymir.claremont.edu
Charles Fumuso	cwf@cray.com
Shari Galitzer	shari@gateway.mitre.org
Tony Genovese	
Robert Hagens	hagens@cs.wisc.edu
Alf Hansen	Alf.Hansen@pilot.cs.wisc.edu
Kevin Jordan	kej@udev.cdc.com
Mukesh Kacker	mukesh@sun.com
Jim Knowles	jknowles@trident.arc.nasa.gov
Ruth Lang	rlang@nisc.sri.com
Mike Little	little@ctt.bellcore.com
John Reinart	reinart@cray.com
Ursula Sinkewicz	sinkewic@decvax.dec.com
Michael Stanton	"usermas@lncc.bitnet"
Michael Tharenos	tharenos@jessica.stanford.edu
Linda Winkler	lwinkler@anl.gov
Russ Wright	wright@lbl.gov

3.5 Operational Requirements Area

Director(s):

- Phill Gross: pgross@nri.reston.va.us

Area Summary reported by Phill Gross/CNRI

Network Joint Management

Gene Hastings presented this report on the activities of the Network Joint Management Working Group. The Working Group followed up on the issue of "unanticipated routing". When things take unanticipated paths, operators need to know about it. Gene and Sue Hares are soliciting anecdotes about pathological cases to get a better understanding of how to handle such exceptions.

Dale Johnson presented some ideas from his Trouble Ticket Wishlist paper and BOF. A mailing list is being creating for the ongoing discussion. Of particular interest was a "quick straw" poll of attendees on how many people have any kind of trouble ticketing system. Ten people said they had some sort of on-line system, whether this happened to be a big ascii file, a variety of mail or an actual on-line database. Four of these were ascii only (which some felt was a virtue because it was printable). Five people had paper-only trouble systems and other people kept it all in their heads. Follow-up discussion dealt with the value of having a database-based trouble ticket system, which might be coupled to other things like configuration and user base. It may be that the internet has gotten very stable, which means that there aren't that many routine problems that bubble up to wide area operators. It also means that what problems do exist are the weird ones which would be exceptions to any trouble ticket mechanism in first place.

It is very important for operators to have ways either to get or correlate information in a timely manner. Many operators make heavy use of "Whois" (subject to connection delays and timely updates of individual entries).

Other issues discussed in NJM were:

- How to deal with international operations. Encouraging amount of activity between FEPG and other concerned parties. Milo Medin pointed out that if you are supporting worldwide operations you must have twenty-four hour on-call or twenty-four hour operations because your midnight may be someone else's day shift.

- What is the operational impact of low-speed or low-end connections?

  - How are you providing these connections either for small organizations, or for individuals?

There was interest in having discussions about new products. This is an interesting issue and it would presumably be very useful to operators to trade experiences with various vendor products. However, this would have to be done very carefully to be fair.

Network Status Reports

Network Status reporting continues to be a standard feature of IETF. A listing of all reports and the slides for the briefings are contained in another section of these Proceedings. However, three reports deserve special mention.

Elise Gerich (MERIT) and Paul Bosco (IBM) reported on the deployment of the T3 NSFnet. (Slides included in these Proceedings) This is something that will be of increasing interest to the entire IETF, and Paul has tentatively volunteered to make a presentation to the plenary by Atlanta (July 1991) or perhaps Sante Fe (Nov 1991). We were very pleased to have Michael Stanton, Associate Professor at PUC-Rio Departamento De Informatica, and Technical Advisor to the Brazilian Research Network, give a very detailed overview of national networking activities in Brazil. (Slides included in these Proceedings) Bernhard Stockman gave a nice report on NORDUnet and European activities.

These last two talks again demonstrate the broadening out of interest and activities in the IETF to a more international scope.

Operational Statistics WG

Operational Statistics had great fun. We began with a very simple model. We hoped to define a common storage format, some common collection tools and some common presentation tools. Then, we would probably use FTP to move files around. After some discussion, we decided it was much more interesting to look at a client/server relationship. Under this model, NOCs could store data in any format they wished locally, but would exchange the data using a server/client relationship in a common exchange format, It gets interesting on the presentation end where you could actually build something like an Xgadget that is wired to the client. Presumably, it could reach out and ask for certain information and bring up a nice Xdisplay on a local terminal in realtime. We began talking about things like a common API for the operational statistics client so that vendors and other developers could create Xtools that would interact directly with these kinds of clients. Of course, it doesn't have to be in real-time. You could always use the simple model that we originally envisioned – that is, query for information and store it locally and present it later.

The Working Group also attempted to define some very generic type of Operational reports. The group felt there needed to be three types of reports for at least three levels of detail – overview (e.g., for management), network engineering and planning, and for realtime troubleshooting. These could roughly be categorized by time granularity. In the first area, there should be a generic monthly report similar to what we tend to see in the Network Status presentations at the IETF. These were called the "McDonald's" reports (e.g. "10,000 packets sold this month"). There is also a need for reporting at roughly daily granularity that would be useful for engineering and network design. That was one of the other major motivations of this activity – to give some common basis throughout the internet for network design. Yet another motivation is trouble shooting and problem resolution, and so you need some sort of instantaneous reporting. The third area received less attention at this meeting.

I think the first two types of "operations reports" were a little clearer to most of the Working Group attendees. Whether we can actually respond in real-time for resolving problems in real-time was largely left for further discussion. That may be more a topic for SNMP.

Late breaking news from a after dinner session was that some folks were proposing ways in which SNMP could be used as the data exchange protocol as opposed to inventing something entirely new.

User Connectivity and Trouble Ticket BOF

UCP and a BOF on Trouble tickets met this week. Each of these groups has a different slant on resolving user connectivity problems. To some extent, it's a difference in scale. UCP is looking at a larger and broader aspect of the problem, while the TTW BOF was focusing more on intra-NOC solutions. My understanding is that these two efforts are reasonably well-coordinated. Both groups have exchanged views, and documents will soon be available for the RFC information track.

"THE PLAN"

1. "DESIGN" THREE GENERIC OPERATIONAL REPORTS TO START WITH.
   - MONTHLY ["MCDONALD'S" REPORT]
   - DAILY [ENGINEERING/NOT DESIGN REPORT]
   - INSTANTANEOUS [PROBLEM RESOLUTION REPORT]

   (WITHIN EXTENSIBLE MODEL, TO ALLOW FUTURE DEVELOPMENT)
   (THIS INCLUDES DEFINING PRESENTATION FORMAT!)
   (START WITH MONTHLY REPORT)

2. CHOOSE METRICS & PARAMETERS (FOR EACH OF THE THREE REPORTS).
3. DEFINE EXCHANGE FORMAT FOR METRICS & DATA EXCHANGE PROTOCOL.
4. CONSIDER PUBLIC DOMAIN TOOLS & PILOT PROJECTS.

MONTHLY REPORT

(METRICS & PRESENTATION FORMAT)

1. NETWORK VOLUME ("MCDONALD'S" REPORT)
2. UTILIZATION (e.g., BUSIEST INTERNAL LINKS)
3. BIGGEST CUSTOMERS (BUSIEST EXTERNAL LINKS)
4. WHAT IS THE TRAFFIC PROFILE (e.g., BY PROTOCOL)
5. CUSTOMER AVAILABILITY (e.g., OUTAGE REPORTS?)

TODAY

→ FINALIZE "MONTHLY REPORT"

EMAIL

→ OTHER ISSUES TOWARD PROTOCOL DESIGN

CURRENT MEETING REPORT

Reported by Dale Johnson/Merit

TTW BOF Minutes

The "NOC Internal Trouble Ticket System Functional Specification Wishlist BOF" met for two full sessions, and also joined UCP and NJM for joint discussions. The group is refining a current internet-draft document describing how NOC internal trouble ticket systems can be integrated with other NOC tools, and what new functionality they could have to make NOC internal operations more reliable and efficient.

The document will be revised for submission as an Informational RFC within a few weeks, to help give direction to several groups currently writing trouble ticket systems.

The group proposed several additions to functionality, and added specific examples and presentation suggestions for the draft. There was quite a bit of discussion about how this proposal could integrate with the current UCP Working Group proposal for an external, inter-NOC national trouble ticket tracking system.

The BOF will have fulfilled its initial purpose within a few weeks when final discussions take place on the mail list and the document is submitted for publication. Trouble ticket ideas will probably continue to be discussed within the UCP and NJM Working Groups, and an informal search has begun to find or develop trouble ticket systems suitable for use in the regionals and campuses. Any such systems found will be proposed to the "Son of NOC Tools" Working Group.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Eric Carroll	eric@utcs.utoronto.ca
Tom Easterday	tom@cic.net
Dale Finkelson	dmf@westie.unl.edu
Eugene Hastings	hastings@psc.edu
Steven Hunter	hunter@es.net
Dale Johnson	dsj@merit.edu
Dan Jordt	danj@nwnet.net
Darren Kinley	kinley@crim.ca
Walter Lazear	lazear@gateway.mitre.org
Mark Leon	leon@nsipo.arc.nasa.gov
Daniel Long	long@nic.near.net
Dan McKernan	mckernan@nsipa.nasa.gov
Milo Medin	medin@nsipo.nasa.gov
Lynn Monsanto	monsanto@eng.sun.com
Donald Morris	morris@ucar.edu
David O'Leary	oleary@sura.net

Rehmi Post	rehmi@ftp.com
Robert Reschly	reschly@brl.mil
Ron Roberts	roberts@jessica.stanford.edu
Karen Roubicek	roubicek@bbn.com
Daisy Shen	daisy@watson.ibm.com
Jim Sheridan	jsherida@ibm.com
Dana Sitzler	dds@merit.edu
Mike Spengler	mks@msc.edu
Bernhard Stockman	bygg@sunet.se
Kannan Varadhan	kannan@oar.net
Ross Veach	rrv@uiuc.edu
Shujiuan Wang	swang@ibm.com
Carol Ward	cward@spot.colorado.edu
John Wobus	jmwobus@suvm.acs.syr.edu
C. Philip Wood	cpw@lanl.gov

3.5.1 Benchmarking Methodology (bmwg)

Charter

Chair(s):

Scott Bradner, `sob@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 have 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:

Done — Issue a document that provides a common set of definitions for performance criteria, such as latency and throughput.

Done — 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.

3.5.2 Network Joint Management (njm)

Charter

Chair(s):

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 efforts 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 Robert J. Reschly, Jr./BRL

NJM Minutes

Agenda

Old Business

- Unexpected routing.
  - Reports
  - Operational Impact
  - Action
    - Is there anything which should be done?
    - Is there anything which can be done?
- Other old issues - Communities?

New business

- Dale Johnson on trouble tickets.

Roundtable on current and expected issues

- Effects of development of Internet
  - Scaling
  - Speed
  - "Low budget" connections, users?
  - International network coordination and mgt., etc.

After a brief review of the function of the NJM, there was another call for "unexpected routing" anecdotes. The University of Delaware to DuPont Delaware via Ithaca, NY and Reston, VA, and WestNET's 16 hops across town routes were cited as examples. Also cited was the TWB routing problem due to that router being connected directly to the campus.

Others mentioned examples which were found to be a result of MILNET problems, and one situation involving Argonne. All were understood and have been or are being corrected.

The subject of diagnosing routing problems came up. Traceroute, especially third-party traceroute where available, still seems to be the most heavily used tool.

Tony Hain of ESNET informed those present of the community name for ESNET's routers. This is strictly for use by other midlevel network operators in the performance of their tasks. Others with a requirement to access these routers should contact Tony. NSI is considering making it's community name available as well.

Dale Johnson briefly outlined this week's discussions concerning NOC trouble ticket systems. He has a draft (draft-ietf-ucp-tt-00.txt), inspired by the UCP Working Group document (draft-ietf-ucp-connectivity-00.txt). He feels that their focus on accountability to end user problem reports and single NOC operations is not totally suitable for his purposes. Dale is more concerned with inter-NOC network oriented operations. Worth noting is that the TT discussions revealed a desire to make this more universally useful – i.e., by central site staff as well as NOC staff. Dale will be publishing an updated document in a few weeks. When questioned about whether any systems were going to be proposed, he responded affirmatively. As a point of information, Gene Hastings stated he felt the real goal of the the UCP paper was the establishment of an inter-NOC transaction processing system for handling the passing of problem reports between NOCs.

MERIT currently runs an IBM mainframe product, but is moving towards a UNIX based TT system they may develop locally. IBM Yorktown is working on xgmon; Tim Salo at MSC is funded to work on a UNIX implementation; and Sun Microsystems is working on one as well. Word on developments will be sent to the Trouble Ticket Requirements mailing list `<noc-tt-req@merit.edu>` (-request for administrivia) as it becomes available.

There was quite a bit of talk about the pros and cons of basing a TT system on top of a DBMS. It is very easy to expend man-years of effort in the design and integration of a DB based system – time many organizations simply do not have. A suggestion that we encourage some company to produce and support a TT system was generally well received. It was also observed that in many cases, the integration of a TT system was going to involve some DB customization/interface work in any event.

A poll was taken about current TT operations. 10 sites have some sort of online TT system (4 were ASCII –['sensibly' printable]); 5 were paper systems; and three people reported having no formal TT system in use. Someone noted there were two publicly available systems (are these in NOCTOOLS?).

Conversation then moved on to the desirabilty of having links to other portions of any existing DB – examples involved things like specification of a router filling in configuration information, and mentioning a pair of routers completing link and telco contact information. Again it was noted that this was a bigger win when the "external" components already existed. It was observed that there must be products available which solve similar problems in areas like inventory control, but that they were not necessarily TT oriented. Unfortunately nobody could cite specific systems.

There was a call to formalize an operations track within the IETF. Having this track would reduce internal schedule conflicts, and should attempt to minimize conflicts with User Services as the two have significant overlap.

The group then dove into an extended discussion of the undesirability of referring all problems up toward MERIT. Members very much wanted the ability to contact relevant parties in other regionals directly, but expressed frustration at lack of contact information. Many rely on one or more of the Internet Managers Phonebook, WHOIS, or stabs into the DNS, but these often are only approximate reflections of reality. One proposal was the addition of text/info records incorporating contact phone numbers.

Doug Gale `<dgale@nsf.gov>` is working on an NSF RFP for global user services.. [something about a help server at MERIT – call (800)66-MERIT and ask about the help server].

There was a suggestion to add DNS records for networks as well as hosts (e.g., lookup on 128.63.0.0 – forward and inverse), along with a warning that any records should match networks.txt.

Milo Medin had some comments concerning the new DDN NIC contract. The new contract does not provide for network number assignment or DNS registration among other things. [Later, Steve Wolff told us that DCA and NSF are working together to ensure the continuity of essential services.] More information will be sent to the mailing list as it becomes available.

Kannan Varadhan then touched on his ongoing Telebit NetBlazer testing. He has developed a list of things he wants to discuss with Telebit, and solicits questions from others. The NJM mailing list `<njm@merit.edu>` (-request for administrivia) will host the dialog with Kannan as his testing continues (i.e., post your questions and answers to this list).

The basic NetBlazer is a 386 box running KA9Q, with 2 modem ports for a total cost of `~$3,000.00`. Additional ports are added in 8 port increments. It offers packet driven dialup, and three authentication methods: username/password; callback; and, between boxes, a crypto handshake. NetBlazer does not do TACACS.

The TACACS comment prompted a number of requests for some sort of authentication servers which may (at least optionally) be Internet-wide in scope. Dale Johnson mentioned in passing that MERIT had just deployed one for MICHNET.

Milo then talked briefly about NSI's plan for having a single 800 number for his folks on travel. When called, this number would route to a hunt group of lines local to that area. He also mentioned that it was still possible to assign fixed IP addresses with this and still have routing work (under OSPF if it was a single area – OSPF used best match.).

After the discussion was wrenched back to the Agenda, it was asserted that overall European routing is a disaster, even if internal (i.e., ESNET or NSI European routing appeared to be sensible). Dave O'leary noted that in many cases routing was set based on technical considerations even when they conflicted with policy considerations. SURA continues to take heat on this issue. It was felt that the FEPG/FRICC work would help. The FEPG has developed guidelines which formalize connectivity in accordance with CCIRN recommendations.

At this point Milo insisted that NOCs contemplating international operation absolutely positively must have 24 x 7 NOC operations.

We were told that SPRINT and Cornell (the NSF International connection managers) want to schedule a global BGP, coordination and cutover meeting. The current best guess has this meeting taking place at the July IETF in Atlanta.

Someone wondered if the decisions were unilateral or bilateral. The IEPG is a technically oriented group doing sensible things, but it is not clear the IEPG is in a position to significantly affect the decision process. Their next meeting is in Paris in early May. It was also noted that many of the problems appeared to be intra-European.

We then moved on to a very brief consideration of what connecting hordes of high schools would entail. A quick survey showed three regionals are planning to connect 10 or more high schools in the coming year, and in at least one case, these connections will connect whole districts.

The humor quotient chose that time to take a significant nosedive so we adjourned.

Attendees

Vikas Aggarwal	vikas@JVNC.net
Jordan Becker	becker@nis.ans.net
Tom Easterday	tom@cic.net
Dale Finkelson	dmf@westie.unl.edu
Vince Fuller	vaf@Standford.EDU
Shari Galitzer	shari@gateway.mitre.org
Joseph Golio	golio@cray.com
Eugene Hastings	hastings@psc.edu
Steven Hunter	hunter@es.net
Dale Johnson	dsj@merit.edu
Dan Jordt	danj@nwnet.net
Daniel Long	long@nic.near.net
Milo Medin	medin@nsipo.nasa.gov
Bill Norton	wbn@merit.edu
David O'Leary	oleary@sura.net
Robert Reschly	reschly@brl.mil
Ron Roberts	roberts@jessica.stanford.edu
Roxanne Streeter	streeter@nsipo.nasa.gov
Kannan Varadhan	kannan@oar.net
Ross Veach	rrv@uiuc.edu
Carol Ward	cward@spot.colorado.edu
C. Philip Wood	cpw@lanl.gov

3.5.3 Operational Statistics (opstat)

Charter

Chair(s):

Bernhard Stockman, bygg@sunet.se
Phillip Gross, pgross@nri.reston.va.us

Mailing Lists:

General Discussion:
To Subscribe:

Description of Working Group:

Today there exist a variety of network management tools for the collection and presentation of network statistical data. Different kinds of measurements and presentation techniques makes it hard to compare data between networks. There exists a need to compare these statistical data on a uniform basis to facilitate cooperative management, ease problem isolation and network planning.

The Working Group will try to define a model for network statistics, a minimal set of common metrics, tools for gathering statistical data, a common statistical database storage format and common presentation formats. Collecting tools will store data in a given format later to be retrieved by presentation tools displaying the data in a predefined way.

Goals and Milestones:

Dec 1990 Agreement on a model.

Dec 1990 Survey for most useful and popular metrics.

Dec 1990 Survey for most useful and popular presentation formats.

Dec 1990 Identify similar efforts being performed by other groups.

Dec 1990 Define a common minimal set of metrics.

Mar 1991 Propose a MIB for metrics not already there.

Mar 1991 Define a common storage format to facilitate data sharing.

Mar 1991 Define common presentation formats to make data comparable.

Mar 1991 Develop outline, and make writing assignments for paper (Opstat1) documenting March 91 milestones.

May 1991 Complete paper Opstat1.

May 1991 Possible mid-term meeting to review Opstat1.

May 1991	Submit Opstat1 as Internet Draft.
Jul 1991	Approve paper Opstat1 for submission as RFC; decide standards-track or Informational?
Jul 1991	Define a new collection of tools based on defined metrics, defined storage formats and defined presentation formats.
Jul 1991	Propose old tools to be retrofitted.
Jul 1991	Develop outline and make writing assignments for paper (Opstat2) on new tools and retrofitted tools
Sep 1991	Complete paper Opstat2
Sep 1991	Possible mid-term meeting to review Opstat2
Sep 1991	Submit Opstat2 as Internet-Draft
Dec 1991	Approve paper Opstat2 for submission as RFC; decide standards-track or Informational?

CURRENT MEETING REPORT

Reported by Dan Friedman/BBN

OPSTAT Minutes

The Operational Statistics Working Group (opstat) met for three sessions. The following report summarizes the proceedings. It is organized along the lines of "Accomplishments", "Issues" and "Process" rather than as a sequential narrative. At the request of the Chairs, the Minutes contain proposals to resolve some of the open issues: basically, a (concrete) cut at what we should do next.

1 Summary of Accomplishments

Our main accomplishments were to agree upon objectives for the work and to take some steps towards realizing those objectives. The objectives are

- To define an architecture for providing Internet access to operational statistics for any Regional or the NSFnet.

- To classify the types of information that should be available.

- To develop (or foster the development of) public domain software providing this information. The aim here is to specify a baseline capability that all the Regionals can support with minimal development effort and minimal ongoing effort. (It is hoped that if they can do it with minimal effort, they in fact will.)

Our progress in each of these areas is described next.

1.1. Architecture

We selected a client/server architecture for providing Internet access to operational statistics, as shown in the figure.

This architecture envisions that each NOC will have a server who provides locally collected information in a variety of forms (along the "raw <--> processed" continuum) for clients. High level proposals for the client/server interaction and functionality for the "first release" of the software are discussed later in the minutes.

1.2. Classification of Opstats Information

We identified three classes of reports based upon prospective audiences. They are:

1. Monthly Reports (a.k.a. "Political Reports") aimed at Management.
2. Weekly Reports aimed at Engineering (i.e., planning).
3. Daily Reports aimed at Operations.

1.3. Development Plan

We decided that it was most important and easiest to address the management reports first, and therefore, we spent the most time focusing on them. We arrived at several key areas:

- Offered Load (i.e., traffic at external interfaces).
- Offered Load segmented by "Customer".
- Offered Load segmented protocol/application.
- Resource Utilization (Link/Router).
- Availability.

The first report came to be known as the "McDonald's Report" (N Billion Bytes/Packets Served).

2. Technical Issues

2.1. Client/Server Interaction

The following was proposed for Client/Server Commands. (The initial proposal was put forth by Dan Long of NEARnet.)

Commands:

- Login (with authentication).
- Help – Returns a description of the available data (names, a pointer to a map, gateways, interfaces, and variables) .
- Format – Defines retrieval format.
- Select/Retrieve – Pose a query to server. (This generates a response containing the data.)
- Exit.

Proposed Query Language:

"SQL-like": SELECT `<router interface>` AND `<variable>` FROM `<startdate>` TO `<enddate>` AT `<granularity>` WITH `<conditions-met>`

The authentication issue was considered important as some of the traffic information, i.e., who's talking how much to whom, will be sensitive. We also felt that the "name/map" issue is important for the following reasons: It will be impossible to agree on a naming structure that is universally meaningful. Even if we could agree on such a convention, it will always be most convenient for the local network operators to maintain information using names that are meaningful to them. Therefore, the server should be permitted to deliver results using the internal names but must able to provide file(s) that enable a person to figure out what the names mean.

Notetaker's Proposal:

Maintain the following information in one or more files. Pointers to information are obtained by the Help command.

Router names:

Gives the name of the router as used in the statistics data. Gives a (human-supplied) description of the router's location, e.g., University XYZ, MegaBig International Corporate Headquarters, or some other information that enables an outsider to determine what role the router is playing in the network. This information embodies the knowledge contained in the network operators' heads.

Net Names:

Provides the (internal) names of the networks attached to the routers' external interfaces. (Router names can be internal here since the information in a) provides a mapping). Gives associated IP addresses.

ASCII file containing backbone point-to-point links (using router names to specify end-points). If the link also has an internal name that will be use when providing link information, give this name. Also gives linespeed. Need to think of a way to specify a connection to a public data service. All data provided by the server is given using internal names.

2.2. Contents of Monthly Reports

We had three presentations on the Monthly Reports (see attached slides). (The groups were commended for their pioneering use of the 11PM-2AM time slot.) Members of the groups were:

- Kannan Varadh? (Photocopy blurred here), Eric Carroll, Bill Norton, Vikas Aggarwal.
- Sue Hares, Et. Al. (Sorry, that's all I have on the hardcopy.)
- Charles Carvalho, Ross Veach, David O'Leary.

The following is a synthesis of the presentations and attendant discussions:

2.2.1. The McDonald's report

The main issues here were: whether to provide packets or bytes or both and whether to provide input or output or both.

Notetaker's Opinion:

I was convinced by the argument that, unless something is radically wrong with the network, differences between input and output should be "down in the noise", and the explanations for the differences will be too obscure for a management report. (If the network is really throwing away a large amount of traffic, we'll hear about it well before a management report has to be written.) So I vote for input only in the McDonald's Report. More on bytes vs. packets later.

2.2.2. Offered Load by Customer

There was agreement that this is useful. The main controversy was how customers should be identified in a publicly available report.

Notetaker's Proposal:

We present the cumulative distribution or density function of offered load vs. number of interfaces. That is: Sort the offered load (in decreasing order) by interface. Plot the function F(n), where F(n) is percentage of total traffic offered to the top n interfaces or the function f(n) where f is the percentage of traffic offered by the n'th ranked interface. (An example appears toward the end of the minutes.)

I feel that the cumulative is useful as an overview of how the traffic is distributed among users since it enable you to quickly pick off what fraction of of the traffic comes from what number of "users." (It will be technically and politically difficult to resolve "user" below the level of "interface.") This graph will suggest more detailed explorations to people who have access to customer "names."

2.2.3. Offered Load by Protocol Type and Application

People seemed to agree that this is valuable and that pie charts are a good way to present the information (since there is no "natural" ordering for the elements of the X-axis, a.k.a, "Category Axis" in spreadsheet lingo.) "By protocol" means TCP, UDP etc. "By application" means Telnet, FTP, SMTP etc. It was also pointed out that it is potentially useful to do this both by packets and by bytes since the two profiles could be very different (e.g., FTP typically uses large packets, Telnet small packets etc.)

2.2.4. Resource Utilization

Everyone agreed that the objectives of this report should be to provide some indication of whether the network has congestion and if/where it needs more capacity. There was considerable debate on exactly how often one would have to poll utilization to determine whether there is congestion and also on exactly what summary statistics to present: averages, peaks, peak of peaks, peak of averages, averages of peaks, peaks of averages of peaks..... We seemed to focus more on link utilization than on router utilization, probably for two reasons. It is more difficult to standardize measures of router utilization, and link costs dominate router costs. We kept looking for some underlying "physics" of networks to determine the collection interval. Here's one opinion.

Notetaker's Opinion:

It will be impractical to determine congestion solely from link utilization, since one would have to collect at a very small interval (certainly less than one minute). Therefore, we should use estimate congestion by looking at dropped packet statistics.

We should use link utilization to capture information on network loading. The polling interval must be small enough to be significant with respect to variations in human activity since this is the activity that drives loading in network variation. On the other hand, there is no need to make it smaller than an interval over which excessive delay would noticeably impact productivity. For example, people won't notice congestion if it only occurs for 10 seconds a day.

30 minutes is a good estimate for the time at which people remain in one activity and over which prolonged high delay will affect their productivity. To track 30 minute variations, we need to sample twice as frequently, i.e., every 15 minutes.

2.2.5. Availability

We didn't have much time to get to this. There was discussion of presenting the information "By Customer" (e.g., Customers with Top N Total Outage Times) or just reporting on # outages that last longer than a certain amount of time.

Notetaker's Proposal:

We should omit Availability reports from the first deployment for several reasons. First, we didn't spend enough time to obtain consensus. Second, they can be politically sensitive. Third, outage data can be very tough to process. Think of trying to determine exactly how a network partition affects connectivity between different pairs of end users. It's an "N-Squared" problem. If we do want to address this, we should start with site, router, and external interface outages only, since these are O(N) problems.

3. Development Proposal

The following is a proposal for a "development/deployment" plan that tries to reach a reasonable compromise among functionality, burden on network operations resources, and "time to market." The discussion is segmented into three parts:

1. What information is to be available through the server?
2. What are the collection/storage requirements?
3. What presentation tools should we build?

3.1. Information Base

The goal of the Server piece is to provide access to data in a fairly raw form (to be described next) and should be the first thing we do. Presentation tools that use this as input can be developed in parallel if people want to but we shouldn't put them on the critical path. We will have to provide the collection tools as well (unless every NOC is already collecting enough data to supply the information outlined below.) The capabilities of the "first release" are to support the:

- McDonald's Report.
- Offered Load by Interface Report.
- Offered Load by Application Report.
- Link Utilization Report.

- Congestion Report.

The Availability Report is missing because it is hard to do and (based upon the level of discussion we had) seemed to be of lower priority. In the first release, we provide a server and client that can deliver the following statistics. For N specified days over a rolling three month interval:

- Total Input Packets and Input Octets per day per external interface.
- Total Input Packets and Octets across the network per day per application. (Note that this is NOT per interface.)
- Mean, Standard Deviation, and Peak 15 minute utilization per day per (unidirectional link)
- Peak discard percentages over fifteen minute intervals per link-direction per day.

The Exchange Format between Server and Client should be ASCII-based because this enables people to quickly look at the data to see if it makes sense and because it enables quick, custom data reduction via AWK. (I have found both these capabilities to be useful in my own analyses of network data.) The first Client that we write should simply retrieve the data in the exchange format and write it to disk. Rationale for this Base:

This information supports the reports described below and then some, so that presentation tools development will not be limited to these reports. The three month collection interval is short enough to keep storage requirements under 5 Mbytes but long enough so that one can examine longer term trends by "dumping" the data a few times a year. (These files should be highly compressible, easily 2:1, since they'll contain mainly ASCII numerals, repetitions of the names of entities, and whitespace, colons etc.) The ASCII-based format will enable us to develop interoperable tools more quickly. TBD:

- The exact exchange format (no real opinion here other than that it be ASCII-based).
- The command structure. The prosed format seems to be an excellent starting point.

3.2. Collection/Storage Requirements

Input bytes and packets per external interface must be collected frequently enough to prevent counter overflow. As they are collected, they can be added to running totals for the day. At the end of the day, the daily totals for each external interface are stored.

Input bytes and packets per application over all interfaces frequently enough to prevent overflow. At the end of the day these can be aggregated into daily totals. (I guess you have collect these per external interface but they can be aggregated into a network-wide total as the day goes on.)

Per link interface per 15 minutes: bytes sent, packets sent, packets received. (To get the drop rate, you have to correlate sent and received at the two ends of the link.) At the end of the day, store away the average utilization, the standard deviation, the peak utilization, and the peak drop percentage. Assuming 10 octets per item for storage, I estimate that the necessary 3 month history can be maintained with <5 Mbytes for a network with 100 routers, 500 external interfaces, and 200 links.

3.3. Reports/Presentation Tools

My hunch is that standardization of presentation tools will come about based on who does the work first. (It's hard to argue with decent code that's in place: to wit, the entire TCP/IP phenomenon.) Here are some suggestions (and the reasoning) for what we should do first.

3.3.1. McDonald's Report

For an N day period, graph Total Input Bytes per day. Put the average packet length as a "note" on the graph.

Reason:

Bytes is a better measure of the "useful" load carried by the network, i.e., the information sent around by the applications; packets are really an artifice of the way we do things. As a network manager, I would be interested in the end-user volume of information. By putting the average packet length, one can convert to packet volumes if need by.

For the same reason, I suggest that the next two reports be done in bytes as well. Note that the suggested initial information base will support comparable presentations by packets as well.

3.3.2. Offered Load by Customer Report

Based on total input bytes for an N day period: Graph the distribution (or density function) of total input bytes vs. external interfaces as shown below. The external interfaces should be put in decreasing order of offered load (in bytes).

3.3.3. Offered Load by Application Report

Based upon total input bytes for the N day period, present a pie chart of the distribution by application.

3.3.4. Link Utilization

The objective here is to provide some information on the utilization of the total set of links and on the "worst" link. The input "data" we have to work with comprises two matrices:

A(i,j) = average utilization of link i on day j
P(i,j) = peak (15 minute) utilization of link i on day j.

Define TAVG(A(i)) = time average of A(i,j) (i.e., sum-over-j(A(i,j))/# days).
Define TAVG(P(i)) = time average of P(i,j) (i.e., sum-over-j(P(i,j))/# days).

I suggest that we order links by the TAVG(P(i)) measure, i.e., the "worst" link is the one that has the highest average peak utilization over the period. Graph the following:

A histogram of the collection of A(i,j) values, using 10X-axis, i.e. plot the function F(n) where F(n) = percentage of A(i,j) entries in the `(n-1)*10% -- n*10%` range.

A comparable histogram of the P(i,j).

Histograms are useful for summarizing the data over all links over the entire period and can suggest further explorations. For the "worst link" (as defined above), plot as a function of day, its average utilization for the day and its peak utilization for the day. (Note that the data that we collect supports exploration of these time series for any link.)

Note that the proposed initial information base will support such analyses for any subset of the links.

3.3.5. Congestion

The available data as specified in section is:

- D(i,j) = peak drop rate (during any fifteen minute interval) for link i on day j.
- Plot a histogram of D(i,j). For the "worst" link (as defined above), say link I,
- Plot D(I,j) as a function of j.

4. Presentations

In addition to the groups on the monthly reports, we had presentations from Bill Norton of Merit and Chris Meyers of Wash. U. (see slides). Chris proposed an exchange format. I'm guessing that the document is available on-line if you wish to review it. Bill discussed Merit's OpStats activities for NSFnet. He focused on their presentation tools as well as the way that they internally organize the data (a tree structure of Unix files). One important point made during this discussion is that relational databases are not good for storing OpStats. (Performance is the issue.) This is unfortunate since many commercial DBMSs are relational in nature, and therefore, we cannot leverage their (usually substantial) report facilities. The idea of a "client/server" model grew out of Bill's presentation.

5. Notable and Quotable

We had some discussion of how Network Managers use Management Reports and, therefore, what the reports need to present. One significant observation was that "Political Graphs don't have to make sense". During Sue Hare's presentation of her group's work on the monthly reports, the KISS acronym was re-interpreted as Keep It Simple Sue.

Attendees

Vikas Aggarwal	vikas@JVNC.net
William Barns	barns@gateway.mitre.org

Eric Carroll	eric@utcs.utoronto.ca
Charles Carvalho	charles@sage.acc.com
Robert Collet	/pnrobert.d.collet/o=us.sprint/admd=telemail/c=us/@sprint.com=
Dale Finkelson	dmf@westie.unl.edu
Daniel Friedman	danfriedman@bbn.com
Demi Getschko	"DEMI@FPSP.HEPNET"
Dave Geurs	dgeurs@mot.com
Fred Gray	fred@homer.msfc.nasa.gov
Olafur Gudmundsson	ogud@cs.umd.edu
Steven Hunter	hunter@es.net
Dale Johnson	dsj@merit.edu
Dan Jordt	danj@nwnet.net
Tracy LaQuey Parker	tracy@utexas.edu
Nik Langrind	nik@shiva.com
Walter Lazear	lazear@gateway.mitre.org
Daniel Long	long@nic.near.net
Gary Malkin	gmalkin@ftp.com
Lynn Monsanto	monsanto@eng.sun.com
Donald Morris	morris@ucar.edu
Bill Norton	wbn@merit.edu
David O'Leary	oleary@sura.net
Rehmi Post	rehmi@ftp.com
Joel Replogle	replogle@ncsa.uiuc.edu
Robert Reschly	reschly@brl.mil
Ron Roberts	roberts@jessica.stanford.edu
Manoel Rodrigues	manoel.rodrigues@att.com
Jim Sheridan	jsherida@ibm.com
Brad Solomon	bsolomon@hobbes.msfc.nasa.gov
Mike Spengler	mks@msc.edu
Bob Stewart	rlstewart@eng.xyplex.com
Roxanne Streeter	streeter@nsipo.nasa.gov
Kannan Varadhan	kannan@oar.net
Shujiuan Wang	swang@ibm.com
Carol Ward	cward@spot.colorado.edu
Cathy Wittbrodt	cjw@nersc.gov
Wing Fai Wong	wfwong@malta.sbi.com
Osmund deSouza	desouza@osdpc.ho.att.com

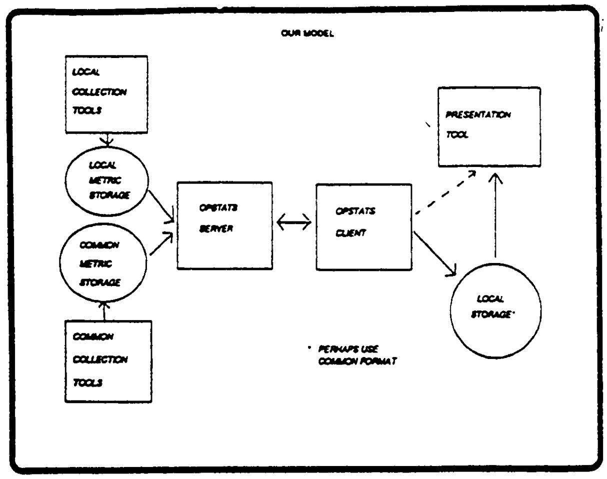
OUR MODEL
LOCAL COLLECTION TOOLS
LOCAL METRIC STORAGE
COMMON METRIC STORAGE
COMMON COLLECTION TOOLS
OPSTATS SERVER
OPSTATS CLIENT
PRESENTATION TOOL
LOCAL STORAGE*
* PERHAPS USE COMMON FORMAT

TOASTER NET

MONTHLY REPORT

JULY 1991

- KANNAN VARADHAN

ERIC CARROLL

BILL NORTON

VIKAS AGGARWAL

... WERE I TO TELL SOMEONE WHO HAS NOT YET SEEN THE BOOK THAT IT CONTAINS FEW WORDS, THAT IT IS FILLED WITH TABLES OF STATISTICS & COLUMNS OF NUMBERS, HE WOULD LOOK UPON THE UNDERTAKING AS A FLOP, EVEN AS INSANITY. BECAUSE WHAT CAN BE DONE WITH HUNDREDS OF PAGES OF STATISTICS.

STANISLAW LEM

"ONE HUMAN MINUTE"

QUESTIONS:

1. NET VOLUME
2. BIGGEST CUSTOMERS
3. USAGE BY SERVICE
4. UTILIZATION
5. CUSTOMER AVAILABILITY

TOASTER LOAD FOR THE MONTH OF JUL '91

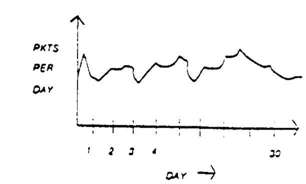

TOTAL: - XXX PKTS / MONTH

OFFERED LOAD TREND (ANNUAL)

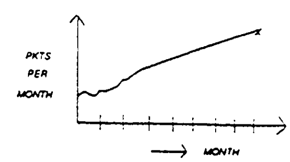

TOASTERNET JULY '91

OFFERED LOAD / CUSTOMER

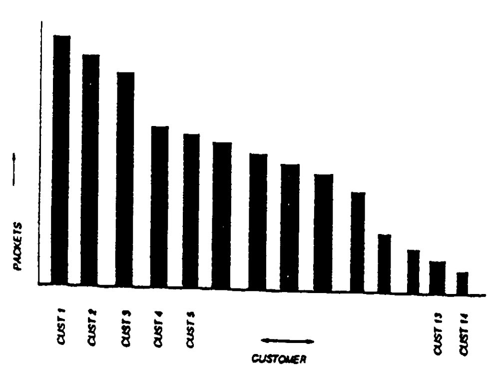

Daily Incoming and Outgoing Traffic—November, 1990

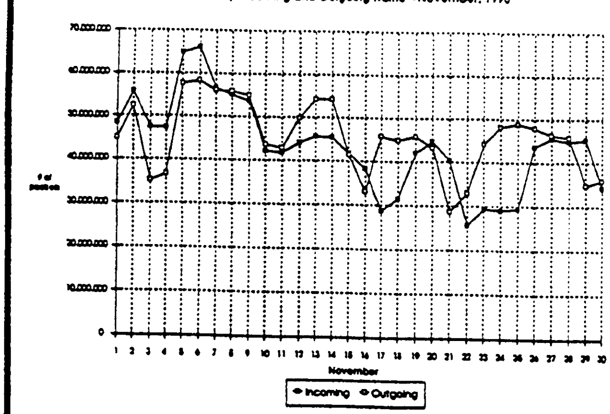

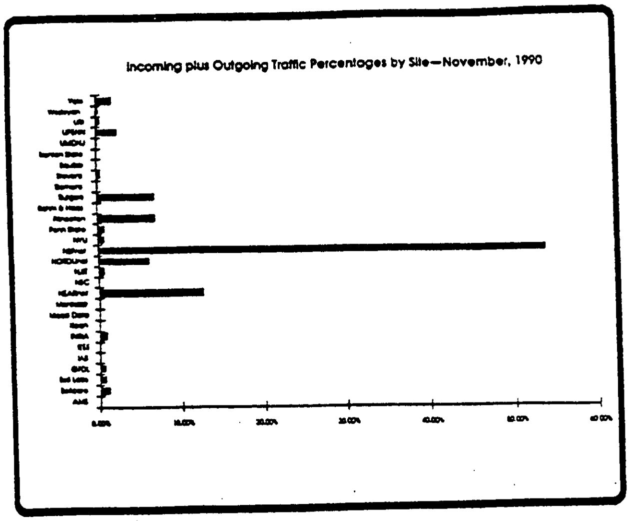
Incoming plus Outgoing Traffic Percentages by Site—November, 1990

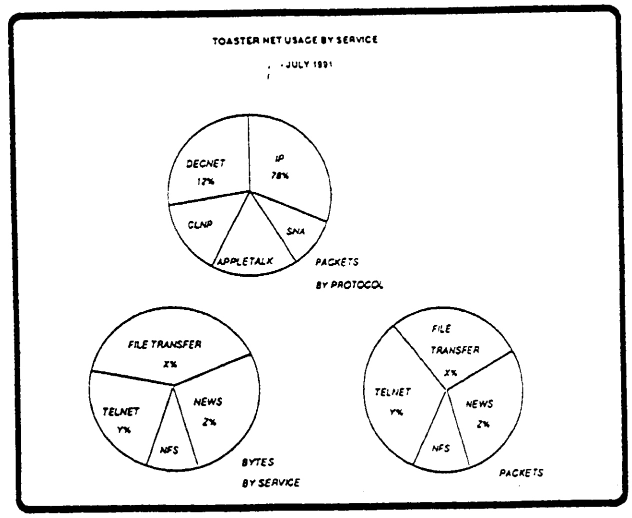
TOASTER NET USAGE BY SERVICE
JULY 1991
DECNET
12%
IP
78%
CLNP
SNA
APPLETALK
PACKETS
BY PROTOCOL
FILE TRANSFER
X%
TELNET
Y%
NFS
NEWS
Z%
BYTES
BY SERVICE
FILE
TRANSFER
X%
TELNET
Y%
NFS
NEWS
Z%
PACKETS

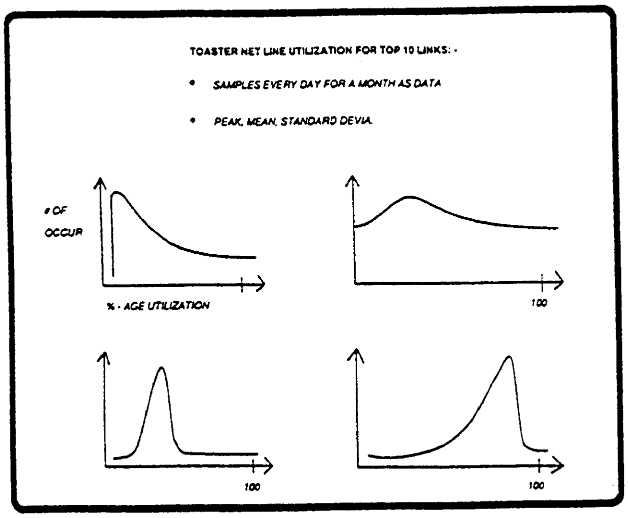
TOASTER NET LINE UTILIZATION FOR TOP 10 LINKS: -
• SAMPLES EVERY DAY FOR A MONTH AS DATA
• PEAK, MEAN, STANDARD DEVIA.
OF
OCCUR
% - AGE UTILIZATION
100
100
100

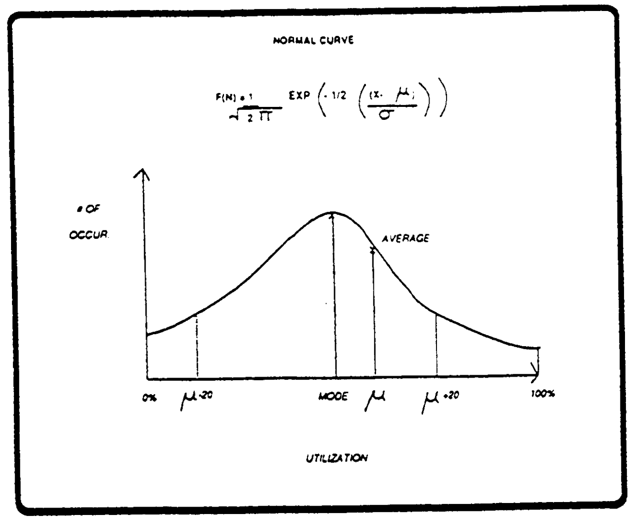
NORMAL CURVE
OF
OCCUR.
AVERAGE
0%
MODE
100%
UTILIZATION

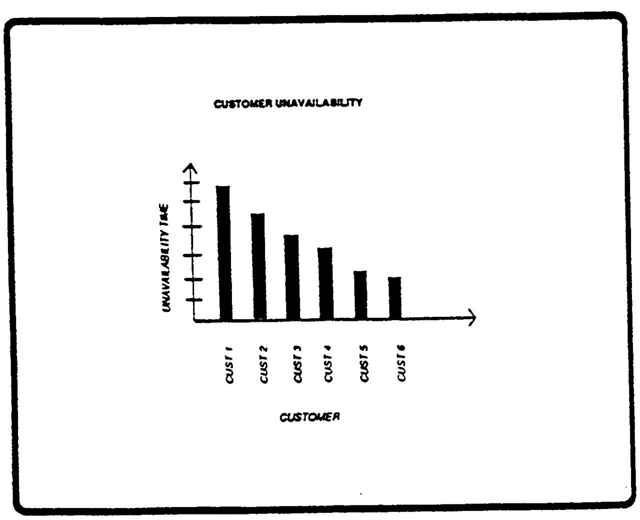
CUSTOMER UNAVAILABILITY
UNAVAILABILITY TIME
CUST 1
CUST 2
CUST 3
CUST 4
CUST 5
CUST 6
CUSTOMER

1. SUE HARES

1.) NET VOLUME

NETWORK COLLECTOR

DAILY PACKETS / BYTES

IN / OUT

- EXTERNAL INTERFACES
- CAVEAT - 2 INTERFACES!

2.) BIGGEST CUSTOMER

- SORT ABOVE INFORMATION BY SITE

3.) WHAT IS IT

DAILY # BYTES / PACKETS

- BY TCP / UDP PORT ISO BY N SELECTOR
- BY PROTOCOL (IP VS. DECNET VS. ISO)
- BY INTERFACE

4.) UTILIZATION

- PACKETS / BYTES   IN / OUT   ? CPU UTILIZATION
- PEAK & AVERAGE   ? CONGESTION
  - PER INTERFACE   ? ERRORS
- POLL TIME 5 MINUTE
- CLIENT CALCULATES UTILIZATION

5.) CUSTOMER AVAILABILITY

- "CLASS 1"   LASTING ≥ 5 MIN

GIVE

DURATION, TIME, DATE, TYPE

CLIENT GENERATES GRAPHS AND

REPORTS -

SERVER - GIVES THE DATA

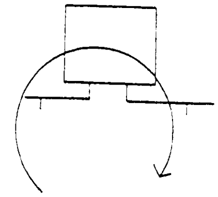

SUE HARES

ET AL.

"clear" data

Charles Carvalho   acc.com
Ross Veach
David O'Leary   sura.net

Question: What is the total network volume?

Select (one of: packets in, packets out, octets in, octets out) starting at 00:00 1-mmm-yy, ending at 00:00 1-nnn-zz, interval 1day, interface (list of external interfaces). This returns an array (items by interval); sum values for each interval, giving total <item> per day across the network.

Question: Where is it going?

Select (one of: packets in, packets out, octets in, octets out) starting at 00:00 1-mmm-yy, ending at 00:00 1-nnn-zz, interval 30days, interface (list of external interfaces). This returns a single value per interface; sort values and take top N (or bottom N).

Question: What is the volume by protocol? traffic profile / offered load

Select (packets by protocol, octets by protocol) starting at 00:00 1-mmm-yy, ending at 00:00 1-nnn-zz, interval 30days, interface (list of external interfaces). This returns an array (protocols by interface); sum values for each interface, giving total <item> per protocol across the network.

Question: What is the (minimum, average, maximum) link utilization?

The average is available from questions 1 and 2: divide octets per day by link bandwidth.

The minimum, maximum values depend on sampling interval, so sampling interval is part of the question: matching in + out (long term vs short term) → loop back, half duplex errors

Select (minimum, maximum) of nn-minute samples of (octets in, octets out), starting at 00:00 1-mmm-yy, ending at 00:00 1-nnn-zz, interval 1day, interface (specific interface or list of interfaces). This returns an array (daily minima or maxima by interface). For single link, show daily numbers; for whole network, compress data for single link into one column. Interesting values: minimum, maximum, average values for month; maximum of daily average?, average of daily maxima?

Question: What is the network availability?

This requires further study. Availability as perceived by a customer cannot be measured directly, as redundant paths mean measuring availability of individual routers or links is insufficient; existence of specific networks in routing table is affected both by network availability and customer activity (i.e., if customer shuts down a router, and stops advertising a network, this should not affect "network availability"

Network Data Collection

Merit Network Inc.

Bill Norton

Merit Computer Network

wbn@merit.edu

Network Data Collection and Storage

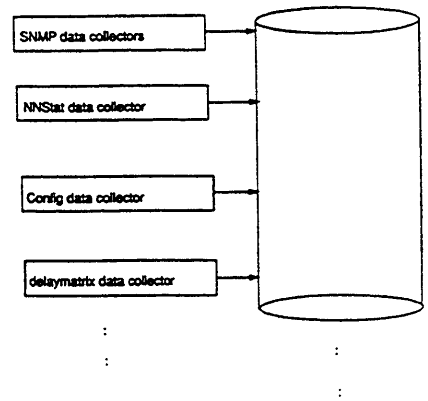

How is the data stored?

Merit Network Statistics Data Collection

Our Environment:

We collect data from multiple networks:

NSFNET
UMNet
Merit
MichNet

We collect multiple types of data:

SNMP
DelayMatrix
NNStat
Off_Net
Configuration (Kernel versions, etc.)

We have collectors for each type of data

We have LOTS of Data

NSFNET: SNMP - 13 Nodes x 7 vars x 96 times/day
8736 SNMP samples per day

Merit: SNMP - 38 Nodes x 2 vars x 24 times/day
1824 samples/day

UMNet: SNMP - 14 Nodes x 7 vars x 96 times/day
9408 samples/day

MichNet: SNMP - 8 nodes x 7 vars x 96 times/day
5376 samples/day

NSFNET: delaymatrix - 13 Nodes x 1 var x 96 times/day
1248 samples/day

NSFNET: NNStat - 40M/day

: : : :

POINT: LOTS OF DATA!!!

Bill Norton
wbn@merit.edu

1

How do we store the data?

FileSystem - Have as much real-time data on-line as possible

Real-Time Data Tree:

/rtdata

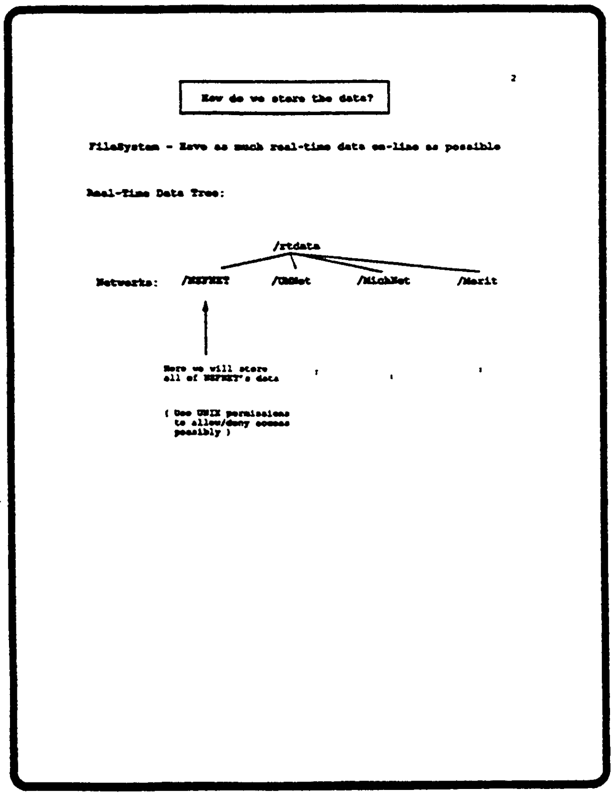
2
How do we store the data?
FileSystem - Have as much real-time data on-line as possible
Real-Time Data Tree:
/rtdata
Networks: /NSFNET /UMNet /MichNet /Merit
Here we will store all of NSFNET's data
(Use UNIX permissions to allow/deny access possibly)

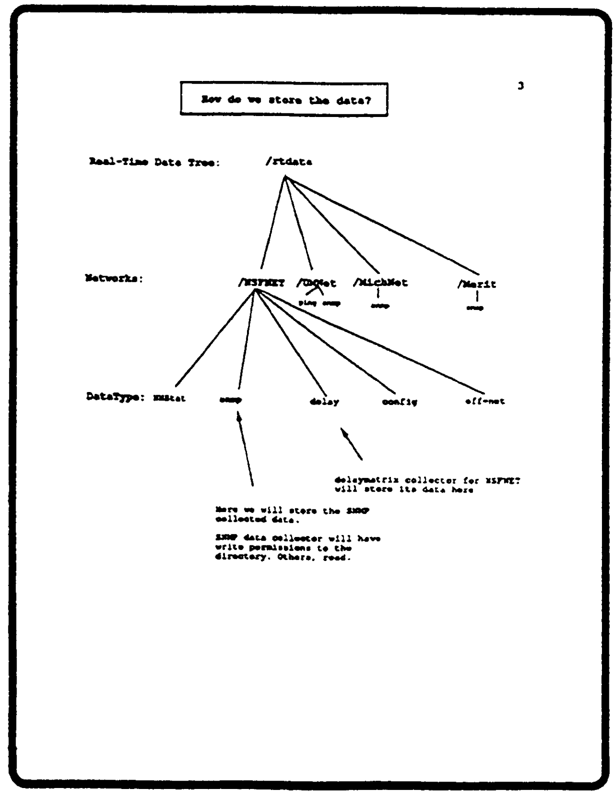
3
How do we store the data?
Real-Time Data Tree: /rtdata
Networks: /NSFNET /UMNet /MichNet /Merit
ping snmp
snmp
snmp
DataType: NNStat snmp delay config off-net
delaymatrix collector for NSFNET will store its data here
Here we will store the SNMP collected data.
SNMP data collector will have write permissions to the directory. Others, read.

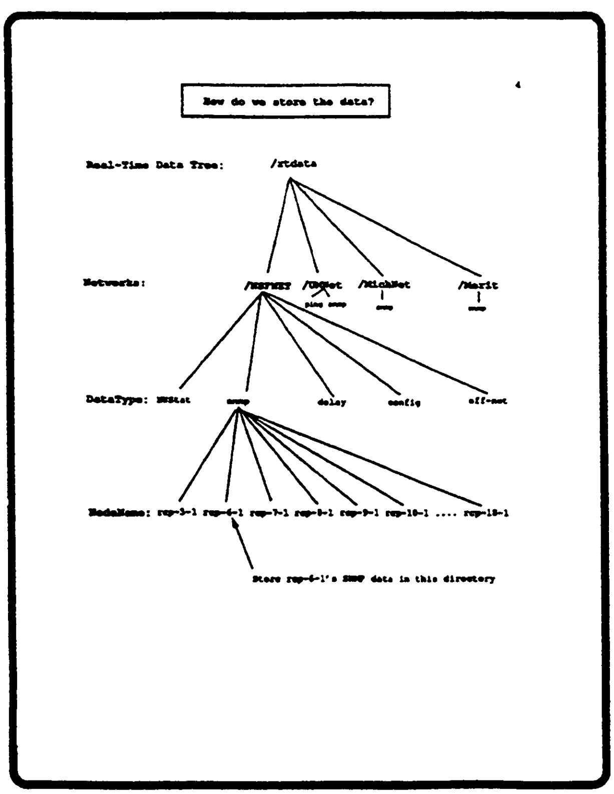
4
How do we store the data?
Real-Time Data Tree: /rtdata
Networks: /NSFNET /UMNet /MichNet /Merit
ping snmp
snmp
snmp
DataType: NNStat snmp delay config off-net
NodeName: rcp-5-1 rcp-6-1 rcp-7-1 rcp-8-1 rcp-9-1 rcp-10-1 .... rcp-18-1
Store rcp-6-1's SNMP data in this directory

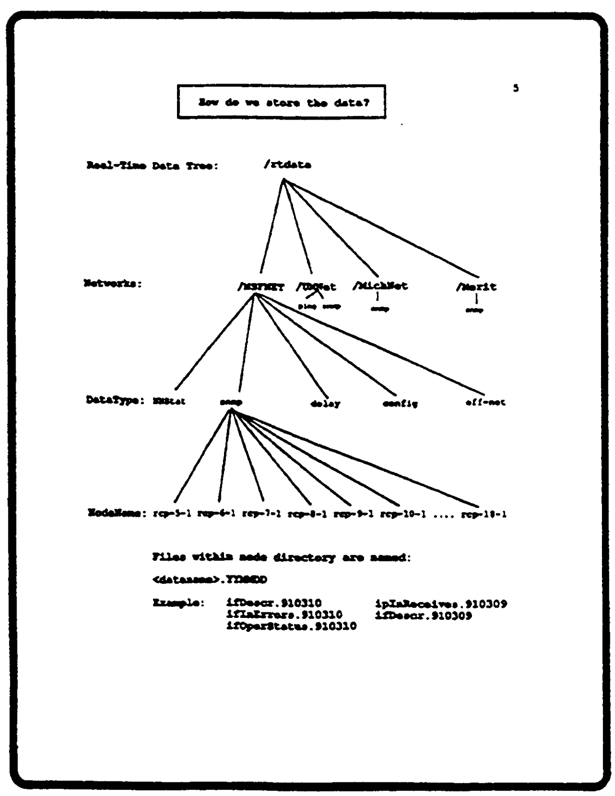
5
How do we store the data?
Real-Time Data Tree: /rtdata
Networks: /NSFNET /UMNet /MichNet /Merit
ping snmp
snmp
snmp
DataType: NNStat snmp delay config off-net
NodeName: rcp-5-1 rcp-6-1 rcp-7-1 rcp-8-1 rcp-9-1 rcp-10-1 .... rcp-18-1
Files within node directory are named:
<dataname>.YYMMDD
Example: ifDescr.910310 ipInReceives.910309
ifInErrors.910310 ifDescr.910309
ifOperStatus.910310

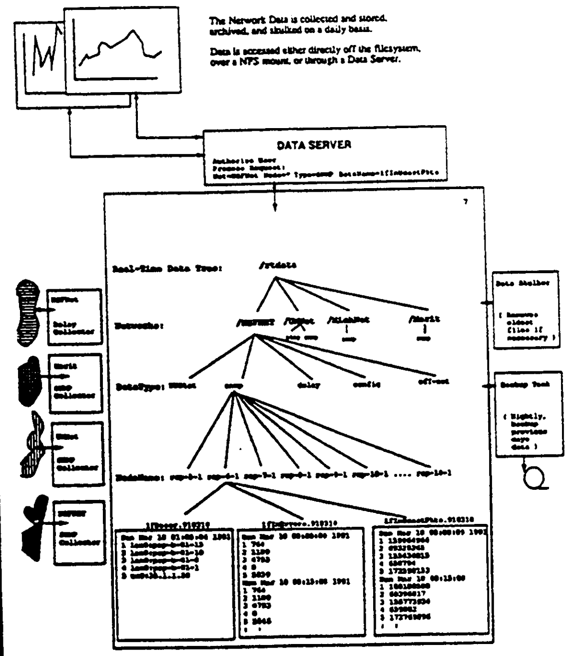

Current Naming Convention

/rtdata/<network>/<datatype>/<node>/<dataname>.<datestr>

```
/rtdata/NSFNET/snmp/rcp-17-1/ifInErrors.910304
/rtdata/NSFNet/delay/rcp-17-1/delaymatrix.910304
/rtdata/UMNet/snmp/Argus-Proteon/ifDescr.910304
/rtdata/Merit/snmp/35.1.1.4/ipInReceives.910304
```

All datestrs and TimeStamps within files are GMT
(World Time - most machines store this way anyway. User programs convert to what they want to see, so make it easy on them too.)

Real Time Data Tree

Bill Norton
wbn@merit.edu

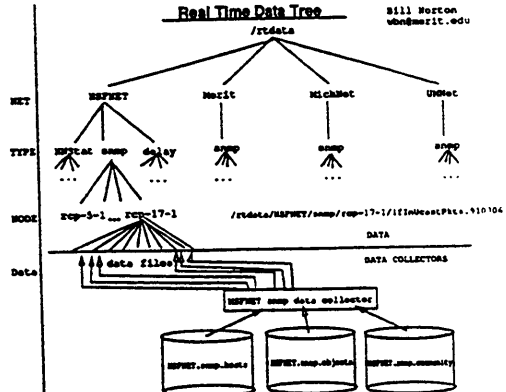

As Data Collectors fill up the hard disk, the data skulker removes old data files.

And after midnight, the backup task tars the previous days' worth of data (using something like tar 'find '*.910306'')

Chris Meyers ①

Proposed Common Data Interchange Format: "Strawman"

The data files are divided into two parts: a raw data file containing a comma-delimited set of data with carriage returns delimiting individual data records, and a descriptor file which contains information about the layout and contents of the data file.

The descriptor file describes what data is being recorded, data ordering information, trivial network topology information (network addresses) and when data samples for the data file were taken. The file format is:

Box: *nickname, description*
Box-Field-1: *tag_number, data_name, data_type, description*
...
Box-Field-*n*: *tag_number, data_name, data_type, description*
Interface-1: *tag_number, <protocol, network_addr>, ..., description*
...
Interface-*n*: *tag_number, <protocol, network_addr>, ..., description*
Interface-Field-1: *tag_number, data_name, data_type, description*
...
Interface-Field-*n*: *tag_number, data_name, data_type, description*
Data-Start-Time: *time*
X-UNIX-Data-Start-Time: *seconds_since_epoch*
X-UNIX-Epoch-Start: *yyyymmddhhmmss* GMT

time	a starting time and date for this dataset in the form yyyymmddhhmmss. *time* is stored in GMT.
tag_number	a unique (for this descriptor file only) integer ≥ 0.
data_name	a string (letters, digits, "-", and "_"). Predefined types are from the MIB: inoctets, outoctets, errors, etc.
data_type	is one of (string, integer, float, boolean, enumerated).
protocol	a string defining the protocol in use. Predefined types include: decnet, ip.
network_addr	is a string identifying the network address of the interface for the specified protocol.
description	is a quoted string identifying the nature and use of the box or interface (hardware, software, maker, purpose, etc). This is freeform.

Chris Meyers ②

A sample descriptor file would be:

Box: ncrc_onch, "ncrc_onch.wustl.edu : cisco brouter : off-campus router"
Box-Field-1: 1, ping_rtt, integer, "Ping Round-Trip-Times from NOC (ms)
Box-Field-2: 2, uptime, integer, "Seconds since router was (re-)booted"
Interface-1: 3, <decnet, 16.1000>, "Ethernet to backbone"
Interface-2: 4, <ip, 128.252.20.1>, "T1 to MIDnet"
Interface-3: 5, <decnet, 16.1000>, <ip, 128.252.123.254>, "Ethernet to Engineering School"
Interface-4: 6, <decnet, 16.1000>, <ip, 128.252.120.254>, <ip, 128.252.135.254>, "Ethernet to Office of the Network Coordinator"
Interface-Field-1: <errors, integer>
Interface-Field-2: <inoctets, integer>
Interface-Field-3: <outoctets, integer>
Interface-Field-4: <available, integer>
Data-Start-Time: 19901225180000 GMT

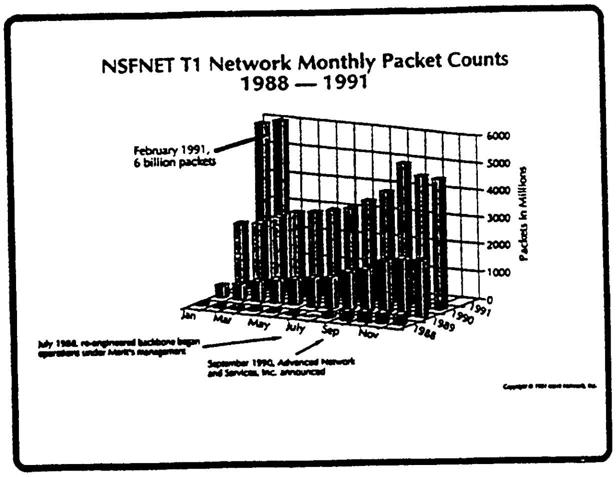

NSFNET T1 Network Monthly Packet Counts
1988 — 1991
February 1991,
6 billion packets
6000
5000
4000
3000
2000
1000
0
Packets in Millions
Jan
Mar
May
July
Sep
Nov
1988
1989
1990
1991
September 1990, Advanced Network
and Services, Inc. announced

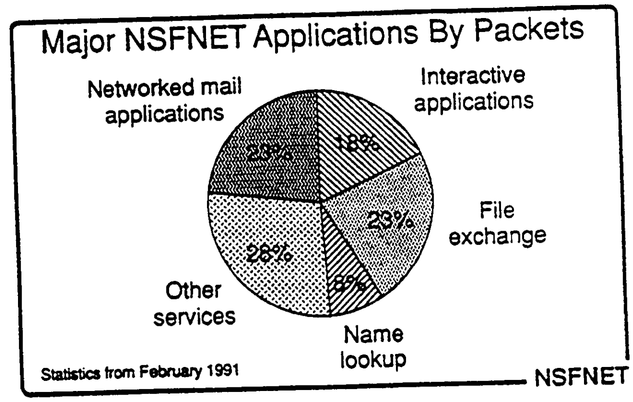

Major NSFNET Applications By Packets
Networked mail
applications
Interactive
applications
23%
18%
23%
File
exchange
28%
8%
Other
services
Name
lookup
Statistics from February 1991
NSFNET

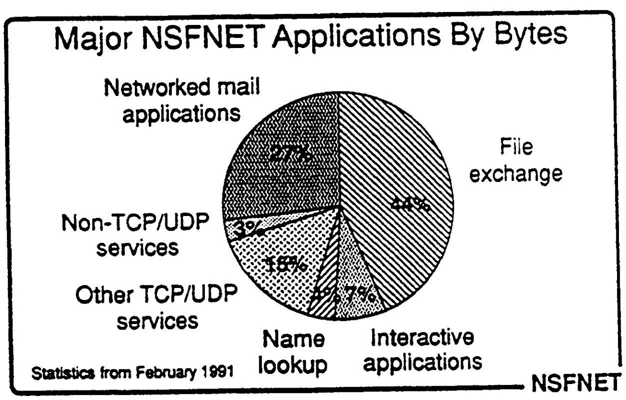

Major NSFNET Applications By Bytes
Networked mail
applications
27%
File
exchange
44%
Non-TCP/UDP
services
3%
15%
Other TCP/UDP
services
4%
7%
Name
lookup
Interactive
applications
Statistics from February 1991
NSFNET

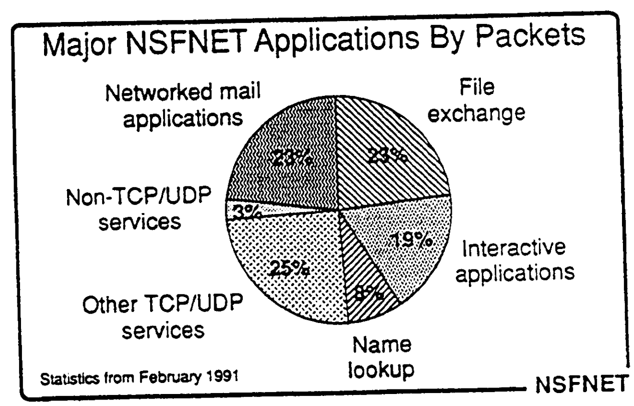

Major NSFNET Applications By Packets
Networked mail
applications
File
exchange
23%
23%
Non-TCP/UDP
services
3%
19%
Interactive
applications
25%
8%
Other TCP/UDP
services
Name
lookup
Statistics from February 1991
NSFNET

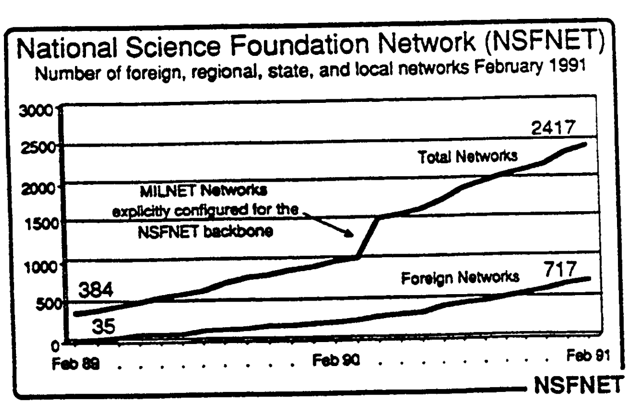

National Science Foundation Network (NSFNET)
Number of foreign, regional, state, and local networks February 1991
3000
2500
2000
1500
1000
500
0
2417
Total Networks
MILNET Networks
explicitly configured for the
NSFNET backbone
717
Foreign Networks
384
35
Feb 89
Feb 90
Feb 91
NSFNET

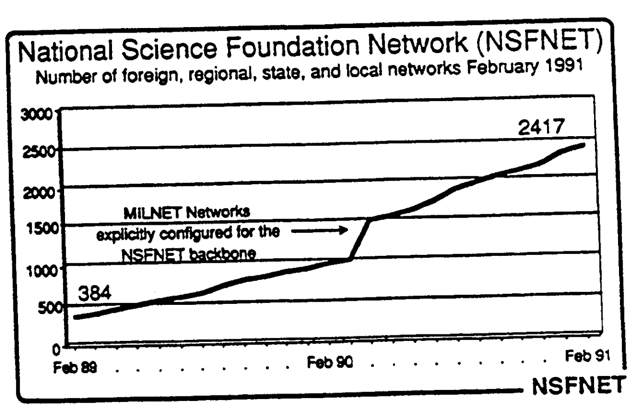

National Science Foundation Network (NSFNET)
Number of foreign, regional, state, and local networks February 1991
3000
2500
2000
1500
1000
500
0
2417
MILNET Networks
explicitly configured for the
NSFNET backbone
384
Feb 89
Feb 90
Feb 91
NSFNET

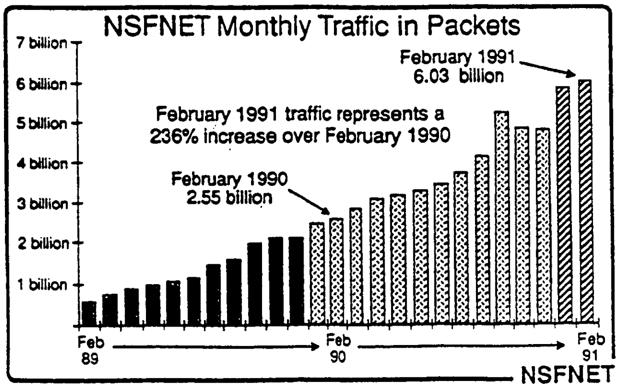
NSFNET Monthly Traffic in Packets
7 billion
6 billion
5 billion
4 billion
3 billion
2 billion
1 billion
February 1991
6.03 billion
February 1991 traffic represents a
236% increase over February 1990
February 1990
2.55 billion
Feb
89
Feb
90
Feb
91
NSFNET

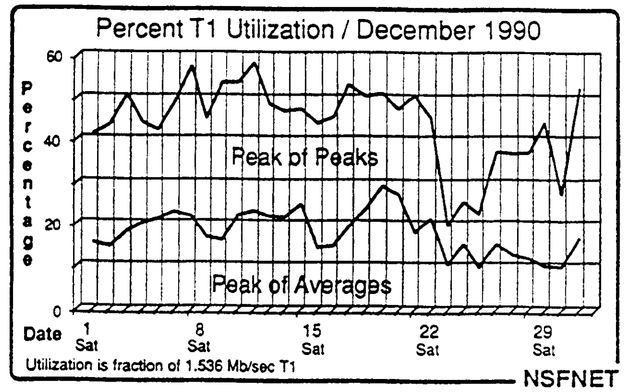
Percent T1 Utilization / December 1990
60
40
20
0
Percentage
Peak of Peaks
Peak of Averages
Date
1
Sat
8
Sat
15
Sat
22
Sat
29
Sat
Utilization is fraction of 1.536 Mb/sec T1
NSFNET

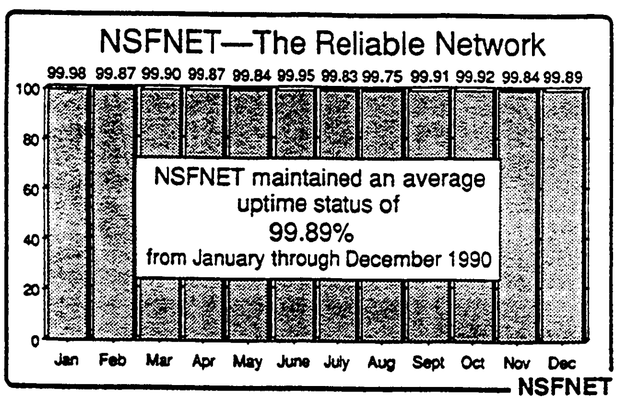
NSFNET—The Reliable Network
99.98 99.87 99.90 99.87 99.84 99.95 99.83 99.75 99.91 99.92 99.84 99.89
100
80
60
40
20
0
NSFNET maintained an average
uptime status of
99.89%
from January through December 1990
Jan Feb Mar Apr May June July Aug Sept Oct Nov Dec
NSFNET

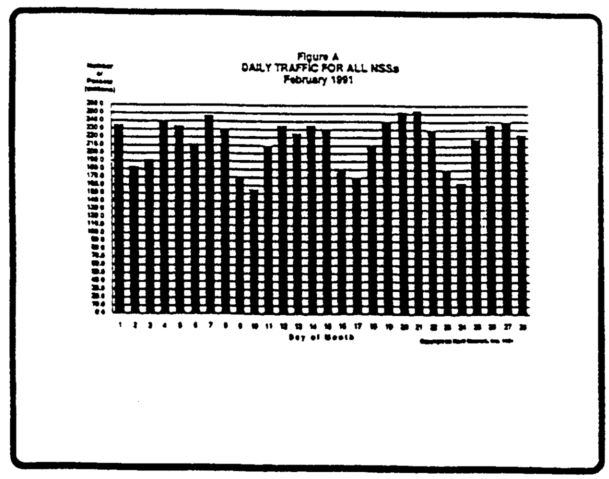
Figure A
DAILY TRAFFIC FOR ALL NSSs
February 1991

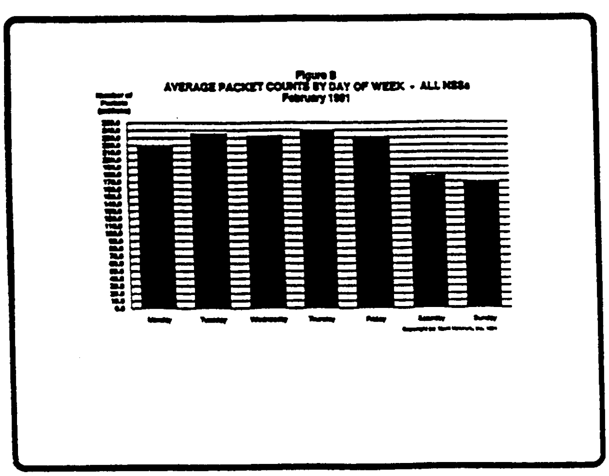
Figure B
AVERAGE PACKET COUNTS BY DAY OF WEEK - ALL NSSs
February 1991

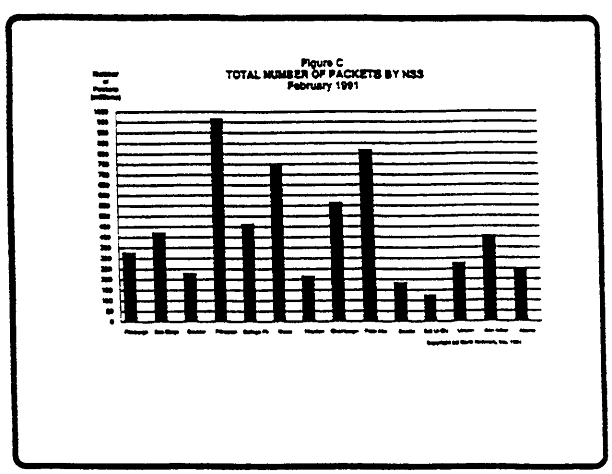
Figure C
TOTAL NUMBER OF PACKETS BY NSS
February 1991

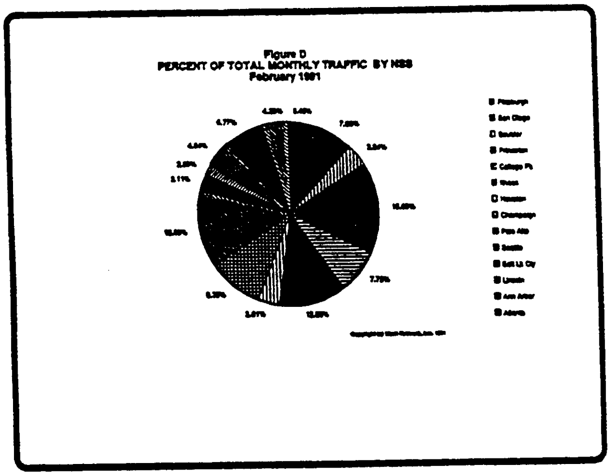
Figure D
PERCENT OF TOTAL MONTHLY TRAFFIC BY NSS
February 1991

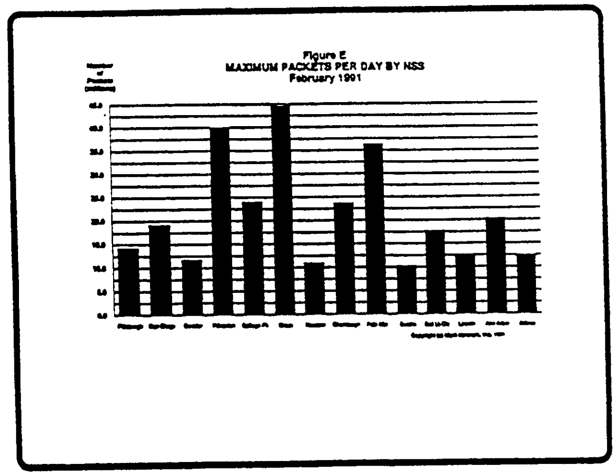
Figure E
MAXIMUM PACKETS PER DAY BY NSS
February 1991

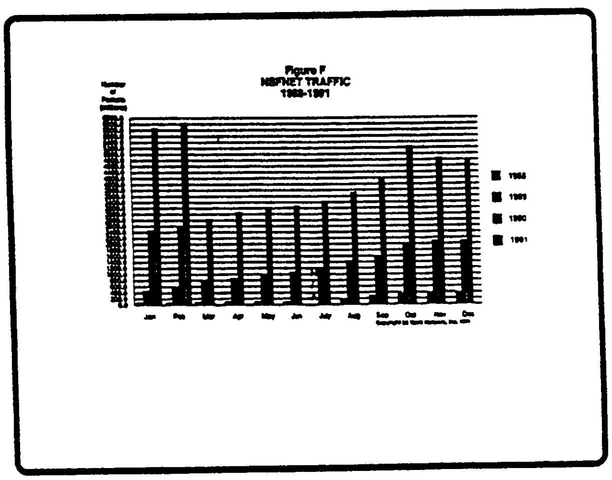
Figure F
NSFNET TRAFFIC

3.5.4 Topology Engineering (tewg)

Charter

Chair(s):
TBD ,

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.

3.5.5 User Connectivity (ucp)

Charter

Chair(s):
Dan Long, long@nic.near.net

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 Dan Long/BBN and Karen Roubicek/BBN

UCP Minutes

Summary

The UCP meeting consisted of a discussion of the UCP Internet Draft. Some of the discussion was to clarify aspects of the draft. The main issue that arose is the obligation for an NSC to accept calls from anyone on any subject. It was agreed that an NSC should be allowed to limit its "liability" by referring callers from outside its customer base or specialty to the "right" NSC. The draft will be ammended to reflect that.

There was also discussion on the issue of the centralized aspects of the UCP plan–how much monitoring is done by the Ticket Support Center and which tickets get tracked by the Ticket Tracking System. There was no consensus on the details but people felt that not all tickets should be tracked and that perhaps suggesting NSC's produce reports on ticket activity would be the most we could "standardize".

There was a brief discussion on the format/mechanism for ticket handoffs but it was acknowledged that we really need some operational experience before suggesting any specifics.

Issues

- Complaints that are dropped between NOC's.
- NOCs that lose tickets.
- Status of problems in design and engineering of networks.
- Statistics on complaints for evaluation.
- Accountability.

Cases

- End host is down.
- MILNET.
- General international connections.
- Anomalous or unexpected routing through experimental networks.
- Telnet options negotiations.
- Vendor software problems.
- General host or applications problems.
- Limitations of low-budget implementations.
- Packet filters.
- Lack of complete problem description.
- Kludge requests.

END HOST IS DOWN - example: user called NEARNET to report that unable to get to andrew.cmu.edu

Two models

1. Nearnet follows through:

```
user----> |nearnet|-----> cmu
    <---  |  nsc  |<----
```

2. Pass off to CMU nsc:

```
user---->nsc---cmu <---> cmu
 |              nsc
 |               |
 \______________/
```

It's rare that a campus would be an NSC because they don't want to track/handle problems outside their campus.

```
(user services)
user -----> campus-----> nearnet
nsc
```

NSC is required to:

- Take ticket regardless of class of problem.
- **Agrees to abide by a core set of rules**.
- Implies responsibility for accepting calls and passing tickets.

Organization can have something outside of its organization that can break rules (like saying "you're not my customer").

Not accept calls <-----> wrong number.
Concern there will be an overload of calls - e.g.: MERIT.

Dana Sitzler: why would a NOC want to be an NSC? What's at the root of the problem? Help users, "support" funding agencies.

Issue of coercion: If I have to take calls, it becomes a funding issue.

Suggest limits on what calls NSC has to take:

- Who must I take calls from?

- What kind of a call must an NSC accept?

```
                                 ---------------
NSC customers: peers            |       |       |
               nsc's            | help  | NSC   |
                                |_______|_______|

Help - acts as filter - redirects people to other NSC's
```

Question of hours or operation: not specified, too variable among organizations

Store information about NSC's in DNS? (eventually); Start with ASCII file

Need to be sensitive to constraints of NSC's

Need to indicate the following for each NSC:

- Customer base
- Scope of expertise

True cost is really too great - need to leverage what exists - pressures regionals to handle more without compensation.

900 number for help? Only real objection seems to be the requirement to accept all calls.

Higher Entity - when NSC's can't get closure, have a frustrated user. But what power does it have?

Text of draft must be revised to recognize:

- Limit scope and customers
- Filter calls

Proposal: NSC's must accept calls from other NSC's but can redirect non-NSC's to other NSC's.

Format for transferring tickets between NSCs (email?).

- TTS supposed to archive completed tickets, have current status (which NSC is holding which ticket)
- Can be an archive of a mailing list
- Authorized NSC's get read access

```
minimum:
To:    new-nsc
cc:    tts
Subj:  Ticket 3076

---------------
---------------
---------------

Limitations of this minimum:  doesn't address who sees what, timers

Only archive (cc:) inter-NSC tickets?
Doesn't address local NOC support issues (what if problem never gets
to TTS?)
```

TSC's are supposed to:

Expedite tickets that aren't making progress, according to timers arbitrate between NSC's act as user ombudsman.

Do all the tickets get reported to TSC? No, not intra-NSC ones. As an alternative, NSC's can do a monthly/weekly report on number of tickets processed, resolved, etc.

How to clarify service requests: Jim Sheridan: some minimal set of requirements for NSC:

- Take trouble tickets.
- Provide reporting on tt's.

Gene Hastings: Rather than requirements, should produce guidelines (at least for publicly-funded organizations) for reporting classes of problems, monthly summaries, etc.

TSC – keeps track of handoffs?

Some service centers have better "clubs" (i.e., leverage).

Classes of calls:

> general info who makes what can't get somewhere who's responsible for... how address mail performance where is a resource unexpected routing is ????? online losing packets protocol X doesn't work application level

Difference between complaint classification vs problem classification (called in as one thing, but turned out to be another) Sheridan: can break down classes into 12 (?) types (hardware, software, connectivity, info, ...)

Reporting recommendations for NSC's - must incorporate into document.

- Jim Sheridan, Gene Hastings, and Dan will draft.
- Is this related to Statistics Working Group?
- Should it be part of monthly report?
- Working Group members will go to OPSTAT meeting and discuss.

To become an NSC, have to agree to rules (define customer base and scope of expertise).

- Accept calls from users in your base
- Follow-up
- Refer to other NSC (redirect)

Should vendors be NSC?

- Sheridan "no"
- O'Leary "yes"

Can publish statistics and put pressure on vendors.

Action Plan

1. Make changes to doc that were discussed.
2. Make recommendation about NSC performance statistics.
3. Maybe someone will implement? write code or procedures?
4. O'Leary will start?

Attendees

Vikas Aggarwal	vikas@JVNC.net
Kathy Atnip	kathy@wugate.wustl.edu
Eugene Hastings	hastings@psc.edu
Steven Hunter	hunter@es.net
Dale Johnson	dsj@merit.edu
Dan Jordt	danj@nwnet.net
Darren Kinley	kinley@crim.ca
Tracy LaQuey Parker	tracy@utexas.edu
Mark Leon	leon@nsipo.arc.nasa.gov
Daniel Long	long@nic.near.net
Lynn Monsanto	monsanto@eng.sun.com
Mark Moody	ccmarkm@umcvmb.missouri.edu
Joel Replogle	replogle@ncsa.uiuc.edu
Ron Roberts	roberts@jessica.stanford.edu
Karen Roubicek	roubicek@bbn.com
Daisy Shen	daisy@watson.ibm.com
Jim Sheridan	jsherida@ibm.com
Dana Sitzler	dds@merit.edu

Mike Spengler	mks@msc.edu
Bernhard Stockman	bygg@sunet.se
Joanie Thompson	joanie@nsipo.nasa.gov

3.6 Routing Area

Director(s):

- Robert Hinden: hinden@bbn.com

See attached slides

ROUTING AREA REPORT

Bob Hinden

March 14, 1991

ROUTING AREA ACTIVITIES AT ST. LOUIS IETF

- Internet Routing Protocol Standardization Criteria Presentation and Discussion
- Presentation on OSPF
- Presentation of BGP
- Border Gateway Protocol W.G. Meetings
- Inter-Domain Policy Routing W.G. Meetings
- IP over Large Public Data Networks W.G. Meetings
- Multicast Extensions to OSPF W.G. Meetings
- Open Shortest Path First W.G. Meetings

BORDER GATEWAY PROTOCOL W.G.
Yakov Rekhter / IBM

- Discussed Extensions to BGP for Inter-AS Routing for Multicast
  - Work was based on Paper by Scott Brim
- Discussed Interactions between BGP and OSPF
- Meet Jointly with the IPLPDN and Discussed Using BGP on Large Public Data Networks
- Discussed How BGP can be used to Heal a Partitioned Autonomous System

INTERDOMAIN POLICY ROUTING W.G.
Martha Steenstrup / BBN

- Tutorial on IDPR on Monday
- Discussed Status of Prototype Implementation and Experiments
  - Setup Policy Routes
  - Changed Topology; Watched Policy Routes Change Correctly
  - All Basic Routing Functions Work
  - Working on Simplifying Configuration Information
  - GATED Implementation in Progress
- Worked on Internet Draft in Detail
- MIB Definition underway

IP OVER LARGE PUBLIC DATA NETWORKS W.G.
George Clapp / Ameritech

- Came to Agreement on Encapsulation of IP Datagrams and Bridged MAC frames on
  - ISDN
  - Frame Relay
- Agreed Upon Means to Identify Protocol Being Encapsulated
- Meet Jointly with the BGP W.G. and Made Progress on the use of BGP on Public Data Networks

OPEN SHORTEST PATH FIRST IGP W.G.
John Moy / Proteon

- Worked on the Reducing the Size of the OSPF MIB
  - Removed One Table and 27 Variables
- Revised MIB will be Submitted for Proposed Status

MULTICAST EXTENSIONS TO OSPF W.G.
Steve Deering / Xerox

- Reviewed Draft Document
- Expect to Have Internet Draft in a Month
- Discussed Interactions between Multicast Inside and Outside an Autonomous System

3.6.1 Border Gateway Protocol (bgp)

Charter

Chair(s):
Yakov Rekhter, `yakov@ibm.com`

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.
Done	Develop a MIB for BGP.
Done	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 Yakov Rekhter/IBM

BGP Minutes

The Border Gateway Protocol Working Group meeting concentrated on the following issues:

- Using BGP as an inter-autonomous system routing protocol for multicast.
- Interaction of BGP with OSPF.
- Using BGP as an inter-autonomous system routing protocol in Large Public Data Networks.
- Repairing partitioned autonomous system with BGP.

Scott Brim submitted a document that proposes several alternatives for using BGP as an inter-autonomous system routing protocol. The Working Group feels that before making any further recommendations it needs more time to study the subject.

The group met jointly with several members of the OSPF Working Group to discuss the issue of interaction between BGP and OSPF. It was agreed that it is very important to produce a document that will unambiguously define such interaction. Several members of both the OSPF and the BGP Working Groups agreed to work on such a document. We expect that the earlier draft of this document will be available before the next IETF.

The BGP Working Group met jointly with the IPLPDN Working Group to discuss how BGP can be used as an inter-autonomous system routing protocol in Large Public Data Networks. The discussion centered around a presentation made by Paul Tsuchiya. Paul agreed to write a document that specifies how BGP should be used in such an environment. As part of this effort, Paul suggested adding new attributes to BGP that would identify the MAC address of the BGP peer. The document that describes this attribute will be posted to the BGP mailing list in the near future.

Dennis Ferguson proposed a mechanism that would allow for repairing a partitioned autonomous system with BGP. This involves the addition of a new attribute. The document that describe this attribute and how it should be used in repairing a partitioned autonomous system will be posted to the BGP mailing list in the near future.

Attendees

Ballard Bare	bare@hprnd.rose.hp.com
Bob Beach	bob@ultra.com
Ronald Broersma	ron@nosc.mil
Richard Colella	colella@osi3.ncsl.nist.gov
Steve Deering	deering@xerox.com

Dino Farinacci	dino@3com.com
Dennis Ferguson	dennis@canet.ca
Jeffrey Fitzgerald	jjf@fibercom.com
Vince Fuller	vaf@Standford.EDU
Dave Geurs	dgeurs@mot.com
Jeffrey Honig	jch@devvax.tn.cornell.edu
Douglas Kerr	dougk@wc.novell.com
Anthony Lauck	lauck@tl.enet.dec.com
Tony Li	tli@cisco.com
April Merrill	
Greg Minshall	minshall@wc.novell.com
Dennis Morris	morrisd@imo-uvax.dca.mil
Gary Mussar	mussar@bnr.ca
Yakov Rekhter	yakov@ibm.com
Stephen Shew	sdshew@bnr.ca
Frank Solensky	solensky@clearpoint.com

3.6.2 IP over Large Public Data Networks (iplpdn)

Charter

Chair(s):

George Clapp, `meritec!clapp@uunet.uu.net`

Mailing Lists:

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

Description of Working Group:

The IP over Large Public Data Networks Working Group (IPLPDN) will specify the operation of the TCP/IP protocol suite over public data networks (PDNs) such as SMDS, ISDN, X.25 PDNs, and Frame Relay. The Working Group will develop and define algorithms for the resolution of IP addresses and for the routing of IP datagrams over large, potentially global, public data networks.

The IP over SMDS Working Group has defined the operation of the Internet protocols when SMDS is used to support relatively small virtual private networks, or Logical IP Subnets (LISs). Issues arising from public and global connectivity were delegated to the IPLPDN Working Group.

The IPLPDN Working Group will also continue the work of the Private Data Network Routing Working Group (pdnrout) on X.25 PDNs. This work will be extended to include call management and the use of the ISDN B channels for the transport of IP datagrams.

Address resolution and routing over Frame Relay will also be discussed.

Goals and Milestones:

Done	Establish priorities and dates of completion for documents.
TBD	Address resolution of Internet addresses to SMDS E.164 addresses, to ISDN E.164 addresses, to X.121 addresses, and to Frame Relay Data Link Connection Identifiers (DLCIs). The algorithm(s) may be defined in either a single or in multiple documents.
TBD	Routing of IP datagrams across very large internets implemented SMDS and on other PDNs.
TBD	Management of ISDN and of X.25 connections and the use of the ISDN B and D channels.

CURRENT MEETING REPORT

Reported by George Clapp/Ameritech

IPLPDN Minutes

Opening Remarks

This was the second meeting of the IP over Large Public Data Networks Working Group and the following was the Agenda of the meeting:

- Wednesday, March 13, 1991
  - Tutorial on Frame Relay by Andy Malis.
  - Discussion of encapsulation and protocol multiplexing over Frame Relay.
  - Tutorial on ISDN (Wayne Heinmiller and Jim Loehndorf).
  - Discussion of encapsulation and protocol multiplexing over ISDN.
- Thursday, March 14, 1991
  - Joint meeting with the BGP Working Group to discuss the use of BGP for routing and address resolution (Paul Tsuchiya).
  - Continued discussion of encapsulation and protocol multiplexing.

Frame Relay Tutorial

Due to airplane troubles, Andy Malis had been unable to present his tutorial on Frame Relay during the plenary the previous evening, so he kindly presented his talk during the first half of Wednesday morning's Working Group session. (Andy also presented this tutorial during the Friday morning plenary, and a copy of his presentation is included in the Proceedings for that session. A postscript version is online for anonymous ftp at "pub/ietf-frame-relay-intro.ps" on ccv1.bbn.com) The presentation was an excellent preparation for the discussion of encapsulation and protocol multiplexing.

Encapsulation and Protocol Multiplexing for Frame Relay

After the tutorial Caralyn Brown presented the following encapsulation and protocol multiplexing scheme for Frame Relay. The proposal is documented in the draft "Multiprotocol Interconnect over Frame Relay Networks", which is available online on ccv1.bbn.com for anonymous ftp as "pub/multiprotocol.txt". (A copy of the viewgraphs of Caralyn's presentation accompanies these Minutes.)

```
+--------------------+
|  LAPD flag (0x7E)  |
+--------------------+
|        DLCI        |
+--------------------+
|        DLCI...     |
+--------------------+
| Format Identifier  |
+--------------------+
|FCSP| Origin_Media  |
+--------------------+
|        ...         |
+--------------------+
|        Info        |
+--------------------+
|        ...         |
+--------------------+
| Frame Ck Sequence  |
+--------------------+
| Frame Ck Sequence  |
+--------------------+
|  LAPD flag (0x7E)  |
+--------------------+
```

FCSP: Frame Check Sequence Preservation

The Format Identifier indicates whether 802.2 and SNAP or a bridged MAC frame follows, and the Origin_Media indicates the type of the bridged MAC frame. If a routed packet is being carried, then the Origin_Media value is set to zero and the 802.2 LLC Type 1 and SNAP headers follow. The Ethertype of the SNAP header is used to indicate the type of the routed packet. This proposal was modified by Charles Carvalho, who proposed that the 802.2/SNAP headers be eliminated by carrying both the bridged MAC frame type and the Ethertype values within a combined Format Identifier/Origin_Media field. There was qualified acceptance of this approach before the group broke for lunch.

ISDN Tutorial and Proposal

Wayne Heinmiller with Jim Loehndorf presented an overview tutorial on ISDN (a copy of the presentation accompanies these Minutes). The presentation led into a talk by Dory Leifer of proposed encapsulation, protocol multiplexing, and fragmentation schemes for circuit and LAPD ISDN. The proposal is documented in the draft "A Subnetwork Control Protocol for ISDN Circuit-Switching", which is available online on terminator.cc.umich.edu as " ftp/isdn" for anonymous ftp. Further discussion was deferred until the next morning.

Joint Meeting with BGP Working Group

Paul Tsuchiya led a joint meeting of the IPLPDN and BGP Working Groups. (A copy of his slides accompany the Minutes.) Paul proposed alternative mechanisms to support a simple and effective way for routers to obtain routing and address resolution information from BGP servers. He also proposed an extension to BGP in which both the next hop IP address and the hardware, or SubNetwork Point of Attachment (SNPA), address would be given to the requesting router.

Members of the BGP group were concerned with potential conflicts between policies of the Autonomous Systems that might be traversed by a packet. This discussion remained unresolved and Paul volunteered to work toward a solution in time for the next meeting.

Encapsulation and Protocol Multiplexing

After the break, IPLPDN met separately from the BGP Working Group and Caralyn Brown presented the modified proposal for encapsulation and protocol multiplexing over Frame Relay. (A copy of the viewgraph of the proposal accompanies the Minutes.)

```
+--------------------+
|  LAPD flag (0x7E)  |
+--------------------+
|        DLCI        |
+--------------------+
|        DLCI...     |
+--------------------+
|  Format ID 1       |
+--------------------+
|  Format ID 2       |
+--------------------+
|         ...        |
+--------------------+
|        Info        |
+--------------------+
|         ...        |
+--------------------+
|  Frame Ck Sequence |
+--------------------+
|  Frame Ck Sequence |
+--------------------+
|  LAPD flag (0x7E)  |
+--------------------+
```

If the value of the Format Identifier is less than 1024 decimal (0x0400) then the field is used to encode the MAC type and the code points are identical to those in internet-draft "Point to Point Protocol Extensions for Bridging". This is shown below.

```
 0                                       1
 0   1   2   3   4   5   6   7   8   9   0   1   2   3   4   5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| 0 | 0 | 0 | 0 | 0 | 0 | F | Z |            MAC TYPE           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

The F bit is used to indicate the presence of the MAC Frame Check Sequence, and the Z bit is used to indicate zero compression.

If the value of the Format Identifier is 1024 decimal (0x0400) then the 802.2 LLC header follows the Format Identifier field, as shown below.

```
 0                                       1
 0   1   2   3   4   5   6   7   8   9   0   1   2   3   4   5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|             DSAP              |             SSAP              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|            Control            |
+---+---+---+---+---+---+---+---+
```

If the value of the Format Identifier is greater than 1024 decimal (0x0400) then the encoded Format Identifier is equivalent to the Digital-Intel-Xerox (DIX) Ethertype, as shown below.

```
 0                                       1
 0   1   2   3   4   5   6   7   8   9   0   1   2   3   4   5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|             Digital-Intel-Xerox (DIX) Ethertype               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Caralyn continued with a discusion of mechanisms for address resolution and proposed that ARP be used to discover the DLCI associated with an IP address. She also proposed an extension to ARP named Inverse ARP to discover the IP address associated with a DLCI. This latter proposal is documented in the draft "Inverse Address Resolution Protocol", which is available online by anonymous ftp on ccv1.bbn.com as "pub/inarp.txt".

After some discussion and modifications, all of these proposals appeared to be acceptable to the group.

IP over ISDN

Dory Leifer continued with a discussion of the following issues:

- Fragmentation - do we need it considering the access network may impose small max frames?
  Resolution: further study
- End-to-end link state for switched access, i.e., XID frames?
  Resolution: no
- ACK mode support or at least include CONTROL field for Q.921 compatibility?
  Resolution: the group accepted the CONTROL field with a pad in the encapsulation scheme. The revised scheme is shown below.

```
            +-----------+
            |   flag    |
+-----------+-----------+
|   DLCI    |  DLCI..   |
+-----------+-----------+
|  Control  |    PAD    |
+-----------+-----------+
| Format ID | Format ID |
+-----------+-----------+
|  info..   |   ...     |
+-----------+-----------+
|   ...     |   ...     |
+-----------+-----------+
|   CRC     |   CRC     |
+-----------+-----------+
|   flag    |
+-----------+
```

- Discovery protocol of Max frame size?
  No resolution
- Code point for in-connection management protocol?
  No resolution

At this point, time ran out and the group adjourned until the next meeting in Atlanta, GA, in July 1991.

Attendees

Karl Auerbach	karl@eng.sun.com
William Babson	bill%penril@uunet.UU.NET
Fred Baker	fbaker@emerald.acc.com

Atul Bansal	bansal@netrix.nac.dec.com
Ron Bhanukitsiri	rbhank@wsl.dec.com
Caralyn Brown	cbrown@wellfleet.com
Howard Brown	brown@ctron.com
Theodore Brunner	tob@thumper.bellcore.com
Lida Carrier	lida@apple.com
Eric Carroll	eric@utcs.utoronto.ca
Charles Carvalho	charles@sage.acc.com
Martina Chan	mchan@mot.com
Cho Chang	chang_c@apollo.hp.com
George Clapp	clapp@uunet.uu.net
Graham Cobb	cobb@marvin.enet.dec.com
Richard Colella	colella@osi3.ncsl.nist.gov
Nabil Damouny	
Tom Easterday	tom@cic.net
Gus Elmer	arch1!ae@att.com
Dennis Ferguson	dennis@canet.ca
Peter Ford	peter@lanl.gov
Daniel Friedman	danfriedman@bbn.com
Charles Gallucci	
Dave Geurs	dgeurs@mot.com
Bret Gorsline	jinx@engin.umich.edu
Martin Gray	mg@spider.co.uk
Olafur Gudmundsson	ogud@cs.umd.edu
Frank Heath	
Kathleen Huber	khuber@bbn.com
B.V. Jagadeesh	bvj@esd.3com.com
Dan Jordt	danj@nwnet.net
Ajay Kachrani	kachrani@regent.enet.dec.com
Manu Kaycee	kaycee@trlian.enet.dec.com
Nam Keung	
Peter Kirstein	kirstein@cs.ucl.ac.uk
Mary Louise Laier	laierl@applelink.apple.com
Michelle Landriault	crm57a@bnr.ca
Joseph Lawrence	jcl@sabre.bellcore.com
Dory Leifer	Dory.Leifer@terminator.cc.umich.edu
Richard Libby	libby@c10sd6.stpaul.ncr.com
Mike Little	little@ctt.bellcore.com
Then Liu	
John LoVerso	loverso@westford.ccur.com
Andrew Malis	malis@bbn.com
Glenn McGregor	ghm@merit.edu
Robert Meierhofer	bobm@c10sd6.stpaul.ncr.com
Linda Melvin	infopath@well.sf.ca.us
Agnes Moran	

Robert Peglar	robp@anubis.network.com
James Philippou	japhilippou@eng.xyplex.com
Rehmi Post	rehmi@ftp.com
Stephanie Price	price@cmc.com
Ron Roberts	roberts@jessica.stanford.edu
Manoel Rodrigues	manoel.rodrigues@att.com
Allyn Romanow	allyn@eng.sun.com
Allan Rubens	acr@merit.edu
Ray Samora	rvs@proteon.com
Paul Selkirk	paul@ftp.com
Marc Sheldon	ms@Germany.eu.net
Daisy Shen	daisy@watson.ibm.com
Stephen Shew	sdshew@bnr.ca
William Townsend	townsend@xylogics.com
Paul Tsuchiya	tsuchiya@bellcore.com
Robert Ullmann	ariel@relay.prime.com
Kannan Varadhan	kannan@oar.net
Francis Waldman	
Wing Fai Wong	wfwong@malta.sbi.com
Chin Yuan	cxyuan@pacbell.com

Multiprotocol Interconnect Over Frame Relay Networks

Caralyn Brown
Wellfleet Communications, Inc.
15 Crosby Drive
Bedford, MA 01730
cbrown@wellfleet.com

CB-03-01 WELLFLEET

Frame Format

LAPD flag (7E hexadecimal)
DLCI*
Format Identifier
FCSP | Origin_Media
Link Protocol Data Unit
LAPD Frame Check Sequence (two octets)
LAPD flag (7E hexadecimal)

Format Identifier
- 802.2 Format 0x00
- Bridged MAC Frame 0x01

FCSP
- No preserved FCS 0
- Preserve FCS 1

Origin Media
- Nonspecific 0x00
- Frame Relay 0x01
- Reserved 0x02
- Ethernet/802.3 0x03
- 802.4 Token Bus 0x04
- 802.5 Token Ring 0x05
- 802.6 MAN 0x06
- To be assigned 0x07-0xff

*DLCIs, as presently defined are two octets. In some networks DLCIs may, optionally be increased to three or four octets. To maintain word boundaries for the data portion of a frame, when using three octet DLCIs, a Frame Relay interface shall pad with an octet of zeros following the three octet DLCI.

CB-03-01 WELLFLEET

Interconnect Issues

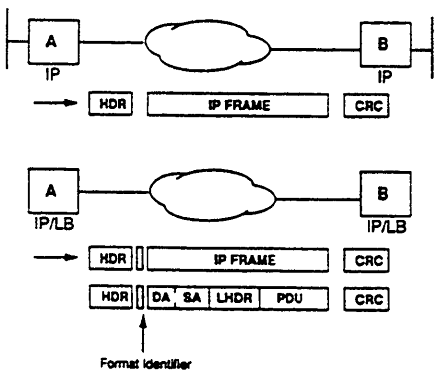

CB-03-01 WELLFLEET

IP Encapsulation for Frame Relay

DLCI (msb)	DLCI (lsb)	
Format Identifier 0x00	FCSP 0	Origin_Media
DSAP 0xAA	SSAP 0xAA	
LLC control 3	(three octet) SNAP	
protocol ID or Org code 0x000000		
Ethertype 0x0800		
IP Packet		

CB-03-01 WELLFLEET

Global Addressing

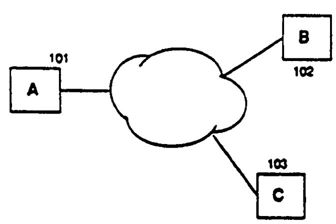

A → B DST DLCI = 102
A → C DST DLCI = 103
C → A DST DLCI = 101
C → B DST DLCI = 102

CB-03-01 WELLFLEET

Local Addressing

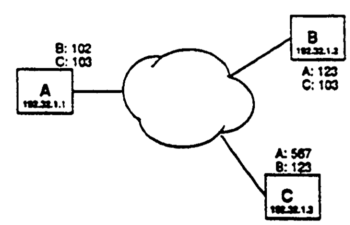

A → B DST DLCI = 102
A → C DST DLCI = 103
B → A DST DLCI = 123
B → C DST DLCI = 103
C → A DST DLCI = 567
C → B DST DLCI = 123

CB-03-01 WELLFLEET

Address Resolution in Locally Addressed Networks

- Station A ARPs for station C
- No source hardware address

request

ar$sha - 0
ar$spa - 192.32.1.1
ar$tha - ?
ar$tpa - 192.32.1.3

- Station C receives request
- Use DLCI from header as source hardware address

request after modification

ar$sha - 123
ars$spa - 19232.1.1
ar$tha - ?
ar$tpa - 192.32.1.3

- Station C swaps target and source source addresses to form response

response

ar$sha - 0
ar$spa - 192.32.1.3
ar$tha - 123
ar$tpa - 192.32.1.1

- Station A receives response
- Use DLCI from header as source hardware address

response after modification

ar$sha - 103
ar$spa - 192.32.1.3
ar$tha - 123
ar$tpa - 192.32.1.1

CB-03-01

WELLFLEET

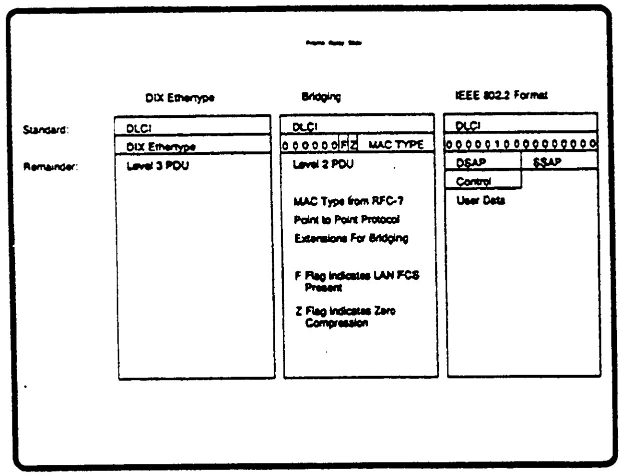

ISDN OVERVIEW

For the Internet Engineering Task Force

March 13, 1991

Wayne Heinmiller (708) 806-8216
Jim Loehndorf (708) 806-8213

Ameritech Services

Outline

Background
- Defining ISDN

Architecture
- Channels, Interfaces, Premises Configurations

Network Services
- Bearer Services
- Supplementary Services
- Teleservices
- User/Application Signaling

Internetworking

DEFINING ISDN

CCITT - International standards
- User- network interfaces

ANSI - National Standards
- Subset of CCITT

Bellcore Technical Requirements - Regional Bell Company Implementations
- Implementation & operational guidelines

Switch Vendors
- Specific implementation details

National ISDN Users Forum
- Implementation agreements for user applications

Access Channels

Access interface is structured in "channels"

B-Channel: 64 kbps clear
- Supports one user device at a time

D-Channel: 16 kbps or 64 kbps
- Supports multiple devices simultaneously
- Used for signaling and (optionally) packet

Miscellaneous Channels
- Embedded Operations Channel, Interface Checksum Channel

Future Channels - Single device at a time
- H0 - 384 kbps (Six times 64 kbps)
- H10 - 1.472 Mbps (T1 minus one DS0)
- H11 - 1.536 Mbps (T1)

D-Channel Details

Uses LAPD for Layer 2 protocol

LAPD (Link Access Procedure for the D-channel) is HDLC derivative

LAPD details
- Layer 3 field carries 128 octets of Q.931 signaling, 260 octets for X.25)
- 2 octet checksum

LAPD Structure

Layer 2 Frame
FLAG
ADDRESS
CONTROL
INFORMATION
CHECK SEQUENCE
FLAG

Flags of "01111110"
- Bit stuffed to prevent spurious flags

Address
- SAPI (0-63) identifies L3 "service" type
  - 0 Signaling, 16 Packet, 63 Management]
- TEI (0-127) identifies physical TE
  - 0-63 User select, 64-126 Ntwk select, 127 for broadcast
- C/R bit identifies frame direction

Control Field
- I Format (for numbered acknowledged info transfer)
- S Format (for flow control, error recovery)
- U Format (for unnumbered info, link control)

Information Field

LAPD Procedures

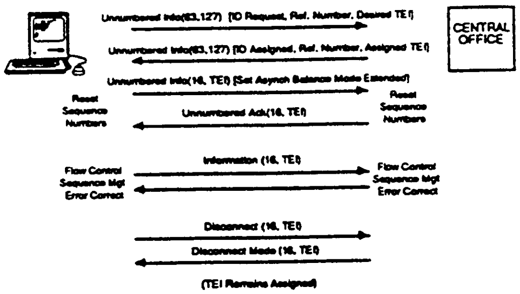

Access Interfaces

Basic Rate (BRI)
- One D-Channel (16 kbps)
- Two B-Channels (64 kbps)
- Miscellaneous channels

Primary Rate (PRI)
- One D-channel (64 kbps)
- 23 B-Channels (64 kbps)
- Miscellaneous channels
- Fills one T1 facility
- Presently can serve only one user device as implemented by vendors

Premises Configurations

Basic Rate
- Supports two physical interfaces
  - U-Interface: 2 wire, supports only one user device, up to 18 kft distance, 160 kbps, 2B1Q (four level) encoding
  - S/T-Interface (bus): 4 to 8 wire, supports multiple user devices, up to 1000 meters, 192 kbps, ASI encoding (binary 1 is no voltage)
- NT1 converts from U-interface to T-interface (NT1 stands for Network Termination 1)

Primary Rate
- U and S/T-Interface: both 4 wire, both 1.544 Mbps, AMI encoding (no difference)
- As presently implemented supports only one user device, but could change in future

Basic Rate Configuration

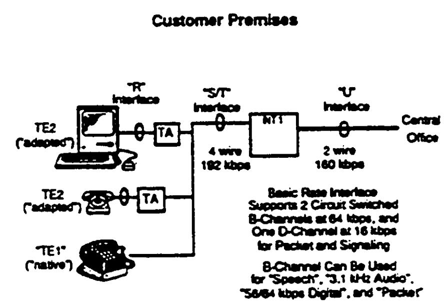

User Devices

Terminal Equipment 1 (TE1)
- Can connect directly to S/T interface
- May connect directly to U interface if it includes built in NT1
- "Speaks native ISDN"
- Examples: voice telephone, PC with ISDN card

Terminal Equipment 2 (TE2)
- Cannot connect to S/T or U interface
- Must be connected to an adapter that connects to S/T or U interface
- Interface to adapter is R interface
- R interface can be RS-232, V.35, etc.
- Examples: PC connected to ISDN terminal adapter using RS-232 interface

Sharing the D-channel

The secret: D-channel bits sent by terminals are echoed by the NT1

Packet vs Signaling
- After closing flag, devices count fill bits (1) before beginning transmission
- Terminals sending signaling can transmit after counting 8 or 9 idle bits
- Terminals sending packets must count 10 or 11 idle bits before beginning transmission
- Terminal must count extra bit if it has transmitted a frame (gives other terminals a chance)

Sharing the D-Channel

(Continued)

Terminal vsTerminal (Layer 2)
- Terminal listens to D-bits echoed from NT1
- If sent 1 but echo is 0, then another terminal is also transmitting
- Terminal detecting collision ceases transmission

Terminal vs Terminal (Layer 3)
- Switch must support layer 3 protocol engine for each terminal
- May need to distinguish between terminals
- Keeps terminal profile, relates to specific terminal during terminal initialization (layer 2 and layer 3 startup)

Ameritech Services — wh & jl 3/13/91

Network Services - Bearer Services

Speech: May be compressed by network(s) during transmission, fidelity not assured

3.1 kHz Audio: May be compressed, but fidelity of typical network transmission is assured. Suitable for analog modems

7 kHz Audio: For higher quality audio transmission, may be used for Group 3 fax. Being renamed Multi-use Bearer Service

Ameritech Services — wh & jl 3/13/91

Network Services - Bearer Services

(Continued)

Circuit Switched 56/64 kbps
- 56 kbps if not over "clear" facilities (eighth bit reserved for network use)
- Rate adaptation for slower data rates carried on channel
- V.110 favored in Europe
  - Structured scheme
- V.120 favored in US
  - HDLC based scheme
- (LAPD+ may be future possibility for rate adaptation)

Ameritech Services — wh & jl 3/13/91

Network Services - Bearer Services

(Continued)

Packet Switched (X.25)
- D-channel: 9600 bps max recognized rate, LAPD layer 2, max virtual circuits 64
- B-channel: 48 kbps max recognized rate, codepoints for 56 and 64 kbps being defined, LAPB layer 2, max virtual circuits 512
- B-channel may be permanently assigned to packet service ("nailed up"), or used on demand (alternating with circuit switched services)

Ameritech Services — wh & jl 3/13/91

Network Services - Bearer Services

(Continued)

Future Bearer Services
- Frame Relay (Switched Virtual Service)
  - Uses modified call control from circuit switched calls (out of band signaling on D-channel) to control virtual circuits on D or B channel.
  - Transport protocol for virtual circuits is subset from LAPD
- H0, H10, H11 Circuit Switched Services
  - 384 kbps, 1.472 Mbps, 1.536 Mbps respectively
- N x 64 Circuit Switched Service
  - From 2 to 24 B-channels aggregated

Ameritech Services — wh & jl 3/13/91

Frame Relay Transport Protocol

Core Aspects of LAPD Frame Structure (LAPD+, Q.922)

Frame Structure

FLAG
Address Byte 1
Address Byte 2
User Data
Checksum Byte 1
Checksum Byte 2
FLAG

User Data is 256 bytes D-channel, 4096 for other channels

Ameritech Services — wh & jl 3/13/91

Call Processing Examples

Circuit Switched
- D-channel signaling (Q.931) is used to establish circuit connection over selected B-channel

Packet Switched
- May use D-channel signaling if B-channel is to be used to carry packet traffic
- Call establishment occurs over selected channel using X.25 protocols

Call Processing Information

Typical Information in Q.931 Message Structure
SETUP Message

- Protocol Discriminator (indicates coding per Q.931)
- Call Reference (identifying number used to identify call in subsequent signaling messages)
- Message Type (indicates SETUP message)
- Sending Complete (indicates no more digits will be sent)
- Bearer Capability (indicates voice, 3.1 kHz audio, or packet)
- Channel ID (indicates channel on which call will be established
- Calling Party Number (originating DN)
- Called Party Number (destination DN)
- Transit Network Selection (identifies IC for interLATA calls)

Originating Circuit Call

Typical Call Control Message Sequences
(Originating a Call)

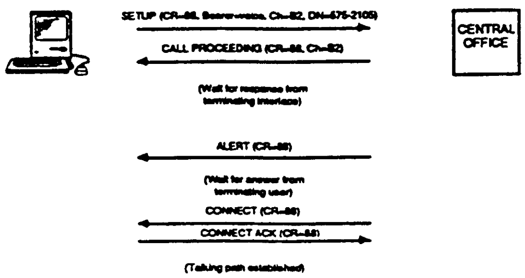

Packet Call Example

Using B-Channel for Originating Packet Call

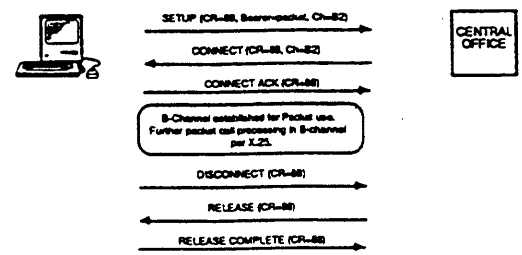

User/Application Signaling

- **Transfer of signaling information between users during call setup may aid user applications**
- **Calling Number - identifies calling party**
- **Called Number - identifies destination if multiple directory numbers are assigned to the interface**
- **Calling and Called Subaddresses - identifies destination beyond network address, examples are LAN address, or PBX extension number**
  - Supports subaddress codings to 20 octets

User/Application Signaling

(Continued)

Low Layer Compatibility
- Identifies rate adaptation and encoding schemes (V.120, 7 bit ASCII, etc.) for circuit switched data calls, format defined but allows user specified codings

High Layer Compatibility
- Identifies application to be served by connection (group 3 fax, Telex, etc.)
- Use favored in some European countries, not in US

User-to-user
- Up to 128 bytes of data, no defined standards, up to bi-lateral agreement between users (passwords, user name, etc.)

Inter-Networking

ISDN users are assigned numbers in E.164 numbering plan
For US (national numbers): N0/1X-NXX-XXXX

To call POTS user (voice), dial telephone number

To call customer on packet network (X.121), send address preceded by numbering plan escape code (0) in X.25 address field of Call Request packet

To call modem user on POTS network, must access modem pool through either ISDN or packet network

Ameritech Services wh & jl 3/13/91

Conclusion

ISDN is powerful, flexible, and can be confusing because of the wide range of options and choices

ISDN can serve a wide range of lower speed (<64 kbps) data needs

ISDN will serve higher speeds (1.5 Mbps, 150 Mbps)

ISDN data services will improve (Frame Relay)

ISDN will be widely available, the start of a national digital (And growing worldwide) network

National ISDN 1 deployment in 1992-93

Ameritech Services wh & jl 3/13/91

BASIC GOAL

- SUPPORT TRANSIT HANDLING OF IP PACKETS
  - ENCAPSULATION
  - DISCOVERY
  - ROUTING

  TODAY'S TOPIC
- OSI PACKETS TOO?
  - PERHAPS ANOTHER GROUP ...

BASIC FUNCTION

- GIVEN DEST. IP ADDRESS
  - FIND SNPA TO FORWARD TO
- ON DEMAND
- PRE-STORED

ONE SOLUTION

(UNACCEPTABLE)

- EVERY ROUTER ON LPDN EXCHANGE BGP INFO. WITH EVERY OTHER

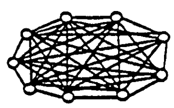

N-1 CONNECTIONS PER ROUTER

BETTER SOLUTION

- ESTABLISH BGP ROUTE SERVERS
  - SIMILAR TO EGP CORE GATEWAYS (HISTORY)

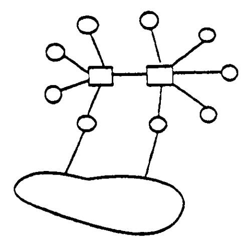

1 CONNECTION PER ROUTER

OPERATION OF BGP

- BGP SERVERS & BGP CLIENTS
  - BOTH RUN STANDARD* BGP
- BGP CLIENT SENDS UPDATE TO BGP SERVER
  - FOR NETWORKS REACHABLE VIA CLIENT
  - NEXT_HOP FIELD CONTAINS BOTH:
    - IP ADDRESS
    - SNPA

    } OF CLIENT
- BGP SERVER ADVERTISES CLIENT'S UPDATE TO ALL CLIENTS
  - NEXT_HOP FIELD UNCHANGED
  - BGP SERVER SHOULD PROBABLY ADD ITS PDN AS TO AS_PATH
- THIS SATISFIES "PRE-STORED" REQUIREMENT

WORK THIS OUT

OPERATION OF BGP (CONT'D)

- CLIENT CAN PEER WITH ONE SERVER AT A TIME, OR MULTIPLE SERVERS

<u>BUT</u> $\boxed{\text{WORK THIS OUT}}$

- CLIENT SHOULD ONLY LISTEN TO UPDATES FROM ONE SERVER AT A TIME
  - SWITCH SERVERS IF PRIMARY GOES DOWN
  - PREVENTS CONFUSION & COMPLEXITY OF DISTINGUISHING POTENTIALLY OUT-OF-SYNC UPDATES FROM MULTIPLE SERVERS.

FINDING BGP SERVER

- CLIENT MUST BE ABLE TO FIND ALIVE BGP SERVER
- OPTIONS:
  - CLIENT STATICALLY CONFIGURED WITH ORDERED LIST OF BGP SERVERS (SNPA & IP ADDRESS)
  - SERVERS ARE MEMBERS OF SMDS MULTICAST GROUP
    - CLIENT ARP's FOR SERVER

OPERATION BETWEEN BGP SERVERS

- EACH BGP SERVER MUST HAVE FULL ROUTING INFO.
  - FULL MEANS UNION OF INFO IN ALL SERVERS
- SERVERS MUST EXCHANGE INFO
  - ALL SERVERS CAN TALK TO EACH OTHER, OR...
  - SERVERS CAN HAVE STILL HIGHER-LEVEL SERVERS
    - SERVER HIERARCHY

ON-DEMAND CLIENTS

- CLIENT IS "DEFAULT" FOR HOSTS BEHIND IT
- ONLY LEARNS APPROPRIATE SNPA WHEN PACKET ARRIVES
- USE ICMP REDIRECT
  - CLIENT SENDS IP PACKET TO SERVER
  - SERVER FORWARDS PACKET TO APPROPRIATE DEST CLIENT
  - SERVER SENDS ICMP REDIRECT TO SOURCE CLIENT
    - PUTS CLIENT IP ADDRESS IN ICMP, NOT SOURCE HOST IP ADDRESS
- OTHER WAYS?
  - SERVER SENDS BGP UPDATE TO CLIENT
- CLIENT DOESN'T WANT TO HEAR BGP UPDATES

ON-DEMAND CLIENTS (CONT'D)

- SERVER CONFIGURED TO SEND NO UPDATES
  - REQUIRES NO BGP CHANGES
  - OR DEFAULTS · REQUIRES "SPECIAL PURPOSE" IMPLEMENTATION
- BGP MECHANISM
  - ~~OPEN MESSAGE INDICATES WHETHER BGP UPDATES DESIRED OR NOT~~
  - ~~RESET BGP CONNECTION TO CHANGE STATES~~
  - ~~OR USE REFRESH MESSAGE?~~

CHANGES TO BGP

- NEED "SWITCH" TO INDICATE WHETHER SERVER/CLIENT FUNCTION DESIRED
  - IF SET, NO THIRD-PARTY ADVERTISING
- MULTIPLE UPDATES, EACH WITH SEPARATE SNPA
  - USE METRICS & DESTINATIONS TO SPLIT TRAFFIC AS DESIRED
    - PRIMARY & BACKUP SNPA's
    - TRAFFIC SPLIT BETWEEN SNPA's
- CURRENTLY BGP DOES NOT CONVEY SNPA
- ADD SNPA FIELD TO NEXT_HOP
- SERVES TWO PURPOSES
  - PREVENTS CLIENTS FROM HAVING TO SEPARATELY DISCOVER SNPA
    - LATENCY, EXTRA PROTOCOL, OVERHEAD
  - ALLOWS BGP ROUTER TO ADVERTISE SNPA's IN ADDITION TO THAT OVER WHICH UPDATES SENT
    - INDEPENDENT OF CLIENT/SERVER BUSINESS

NEW BGP NEXT_HOP ATTRIBUTE

NAME	TYPE CODE	LENGTH	CATEGORY
NEXT_HOP	*4*	*VARIABLE*	*WELL-KNOWN DISCRETIONARY*

FIELD	LENGTH
ADDRESS	*4 OCTETS*
SWITCH	*1 OCTET*
SNPA	*VARIABLE* *LENGTH - 5*

- SNPA ON LSB SIDE OF FIELD, "0" PADDING ON MSB SIDE
  - SYSTEM MUST KNOW SNPA LENGTH (IN BITS)

3.6.3 ISIS for IP Internets (isis)

Charter

Chair(s):
Ross Callon, callon@bigfut.enet.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.

3.6.4 Inter-Domain Policy Routing (idpr)

Charter

Chair(s):
Martha Steenstrup, msteenst@bbn.com

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

Description of Working Group:

The Inter Domain Policy Routing Working Group is chartered to develop an architecture and set of protocols for policy routing among large numbers of arbitrarily interconnected administrative domains.

Goals and Milestones:

Done	Write an architecture document.
Done	Draft Protocol Specification of key elements of the protocol.
Done	Develop a prototype implementation of the protocols.
Ongoing	Gain experience with the prototype in "real networks".
TBD	Develop gated version.
TBD	Add a small set of additional features and submit protocol into IETF standards process.

CURRENT MEETING REPORT

Reported by Martha Steenstrup/BBN

IDPR Minutes

The Inter-Domain Policy Routing (IDPR) Working Group met for four sessions at the March 1991 IETF meeting in St. Louis. As we have many new members at each IETF, we always provide a tutorial as the opening session of the IDPR Working Group.

The second session was devoted to a discussion of the IDPR prototype implementation, concentrating on those parts of the implementation that are complicated and on the configuration information. We also described results of experiments in our test labs and on DARTNET. Our experiments were of the proof-of-concept variety: IDPR does indeed generate routes respecting domain transit policies and adapts routes to changes in connectivity and policy. During the final two sessions, we covered the IDPR protocol specification draft in detail, soliciting comments and suggestions for improvement. We invite all interested parties to read and comment on the Internet Draft. Please send responses to idpr-wg@bbn.com.

Currently, members of the IDPR Working Group are in the process of putting together a MIB, a gated version of IDPR, and a user guide. Once these become available, people interested in experimenting with IDPR will be able to do so with relative ease.

Attendees

Atul Bansal	bansal@netrix.nac.dec.com
Helen Bowns	hbowns@bbn.com
Scott Brim	swb@devvax.tn.cornell.edu
Robert Collet	/pn=robert.d.collet/o=us.sprint/admd=telemail/c=us/@sprint.com
Don Coolidge	coolidge@speaker.wpd.sgi.com
Barbara Denny	denny@sri.com
Dino Farinacci	dino@3com.com
Elizabeth Feinler	
Peter Ford	peter@lanl.gov
Demi Getschko	"DEMI@FPSP.HEPNET"
Dave Geurs	dgeurs@mot.com
Fred Gray	fred@homer.msfc.nasa.gov
Douglas Kerr	dougk@wc.novell.com
Mary Louise Laier	laierl@applelink.apple.com
Tony Li	tli@cisco.com
Mike Little	little@ctt.bellcore.com
Bill Manning	bmanning@houston.sc.ti.com
Mike Marcinkevicz	mdm@calstate.edu
Glenn McGregor	ghm@merit.edu
Carol Melowitz	melowitz@mdcgwy.mdc.com
April Merrill	

Greg Minshall	minshall@wc.novell.com
Dennis Morris	morrisd@imo-uvax.dca.mil
John Moy	jmoy@proteon.com
Andy Nicholson	droid@cray.com
Stephanie Price	price@cmc.com
Manoel Rodrigues	manoel.rodrigues@att.com
Stephen Shew	sdshew@bnr.ca
Frank Solensky	solensky@clearpoint.com
Martha Steenstrup	msteenst@bbn.com
Jonathan Wenocur	jhw@shiva.com

3.6.5 Multicast Extentions to OSPF (mospf)

Charter

Chair(s):
Steve Deering, deering@xerox.com

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/Xerox PARC

MOSPF Minutes

Agenda

1. Review John Moy's draft specification of March 4.
2. Discuss outstanding issues:
   - Inter-area delivery trees
   - Inter-AS delivery trees
   - Multicast group scope
3. What next?
   - Progress of specification
   - Implementations?
   - Testing?

Minutes

Most of the meeting was spent on Agenda Item 1, a page-by-page review of John's latest spec. We agreed to one significant change: to merge the separate concepts of the "SPF cache" and the "forwarding cache", and to refer to the "local group cache" as something other than a "cache", perhaps a "list" or "table". There were also a number of minor or editorial changes suggested, that John will incorporate in the next draft.

Under Agenda Item 2, we discussed (again) the alternatives for inter-area and inter-AS routing. The group agreed to the inter-area scheme proposed by John, which is based on RPF for identifying entry routers, but using FPF within an area. The inter-AS approach, however, remained unclear, pending decisions by the BGP Working Group, which is considering a number of proposals by Scott Brim. Steve and Scott also described a couple of proposals for adding some sort of scope limits to group memberships (orthogonal to the use of TTL for scope of multicast transmissions); possibilities include a scope subfield within an IP multicast address, or a separate scope parameter in the Join Group operation (and in the resulting IGMP message). This remains a topic for future consideration.

For Agenda Item 3, John agreed to incorporate the suggested changes and to flesh out the missing sections of the spec, except for those parts dealing with inter-AS routing, and to make the document available as an internet draft. Scott and Jeff Honig at Cornell are planning to implement the OSPF multicast extensions within "gated" as time and resources permit.

Attendees

William Babson	bill%penril@uunet.UU.NET
Scott Brim	swb@devvax.tn.cornell.edu
Rob Coltun	rcoltun@trantor.umd.edu
Steve Deering	deering@xerox.com
Barbara Denny	denny@sri.com
Kurt Dobbins	dobbins@ctron.com
Dino Farinacci	dino@3com.com
Dennis Ferguson	dennis@canet.ca
Jeffrey Honig	jch@devvax.tn.cornell.edu
Anders Klemets	klemets@cs.cmu.edu
Michelle Landriault	crm57a@bnr.ca
Bill Manning	bmanning@houston.sc.ti.com
Greg Minshall	minshall@wc.novell.com
John Moy	jmoy@proteon.com
Gary Mussar	mussar@bnr.ca
David O'Leary	oleary@sura.net
Steven Sherry	shsherry@eng.xyplex.com
Stephen Shew	sdshew@bnr.ca
Frank Solensky	solensky@clearpoint.com
Paul Tsuchiya	tsuchiya@bellcore.com
Osmund deSouza	desouza@osdpc.ho.att.com

3.6.6 Open Shortest Path First IGP (ospf)

<u>Charter</u>

Chair(s):

Mike Petry, petry@trantor.umd.edu
John Moy, jmoy@proteon.com

Mailing Lists:

General Discussion: ospfigp@trantor.umd.edu
To Subscribe: ospfigp-request@trantor.umd.edu

Description of Working Group:

The OSPF Working Group will develop and field test an SPF-based Internal Gateway Protocol. The specification will be published and written in such a way so as to encourage multiple vendor implementations.

Goals and Milestones:

TBD Design the routing protocol, and write its specification.

TBD Develop multiple implementations, and test against each other.

TBD Obtain performance data for the protocol.

TBD Make changes to the specification (if necessary) and publish the protocol as an RFC.

CURRENT MEETING REPORT

Reported by John Moy/Proteon

OSPF Minutes

The OSPF Working Group met Thursday, March 14th at the St. Louis IETF.

The main topic of discussion was the OSPF MIB. Fred Baker and Rob Coltun had published the OSPF MIB as an Internet Draft, and had received a number of comments, especially from the Network Management Directorate. Their chief comment was that there were too many variables. So the Working Group went through the MIB variable by variable, to see which variables could be combined and which could be discarded. We ended up reducing the 127 original variables to 101.

A number of other things were also discussed:

- Fred Baker pointed out that, for all MIB groups that can be dynamically created/destroyed, there needs to be an extra variable to accomplish deletion. This will add to the size of the MIB.

- Some time was spent in discussing how to do route filtering in OSPF at area boundaries. This can be done through selective non-advertisement of certain area ranges. John Moy pointed out that this is dangerous to do in transit areas. It was decided that since this type of route filtering is not covered in the OSPF spec, it would not be mentioned in the OSPF MIB either.

- Dino Farinacci mentioned a problem he encountered where an OSPF router (call it router X) received an external LSA originated by router Y which listed one of X's IP addresses as the forwarding address. After performing its routing table calculation, router X ended up pointing at itself for the destination. It was agreed the the spec should say that, when performing the routing table calculation, those external LSAs whose forwarding address is one of the router's own addresses should be ignored. [Editor's note: The very existence of such an advertisement probably indicates a problem in the exchange of information between OSPF and other protocols such as RIP/EGP].

- We agreed that all configuration items in the MIB will have read/write access, but that the MIB will say that in the absence of adequate authentication the variables can be implemented as read-only.

There was not enough time at the meeting for examination of the OSPF Trap MIB that Rob Coltun has created. We plan on discussing this document via electronic mail.

Attendees

Fred Baker	fbaker@emerald.acc.com
Rob Coltun	rcoltun@trantor.umd.edu
Dino Farinacci	dino@3com.com
Jeffrey Honig	jch@devvax.tn.cornell.edu
Ann Kerr	
Donald Merritt	don@brl.mil
John Moy	jmoy@proteon.com
Michael Reilly	reilly@pa.dec.com
Glenn Trewitt	trewitt@pa.dec.com

3.7 Security Area

Director(s):

- Steve Crocker: crocker@tis.com

Area Summary reported by Steve Crocker/TIS

Security Track at St. Louis

Monday, March 11

9:30am	Security Area Advisory Group meeting (Stephen Crocker/TIS) SAAG members are required to attend; all IETF attendees are welcome to observe. The status of security related activities and plans for the meetings of the week will be discussed.
1:30pm	Common Authentication Protocol BOF (John Linn/DEC) This is a formative meeting for this group. It will discuss the creation of a standard authentication service interface and the issues surrounding the creation of such a standard, which is not normally pursued within the IETF. Relevant IESG members are expected to be present.
4:00pm	Password Management BOF (Jeff Schiller/MIT) This is a formative meeting for this group. This group is expected to address the problem of sending cleartext passwords across a network. A final decision as to the scope and actual direction of the group will be discussed at this, its first, meeting.

Tuesday, March 12

9:00am	Security Policy (Rich Pethia/CERT) The revised document will be presented for approval. The group will discuss its submission as an Internet standard. Router Requirements (Philip Almquist/Consultant) Members of the SAAG will be present at this meeting to discuss security issues.
1:30pm	Site Security Policy Handbook (Paul Holbrook/CICNet and Joyce Reynolds/ISI) The document will be reviewed with the expectation that it will be submitted as an Internet-Draft as soon as possible after the meeting.

Privacy Enhanced Mail I BOF (James Galvin/TIS)

RFCs 1113, 1114 and 1115 have been revised and are expected to be submitted to the standards process soon after this meeting. This group will discuss some of the more significant changes to the RFCs.

7:00pm Privacy Enhanced Mail II BOF (James Galvin/TIS)

The deployment of Privacy Enhanced Mail will set in place an infrastructure to address the creation and maintenance of certificate-based key management. This will require organizations who issue certificates to sign an Organizational Agreement with RSA DSI. Jim Bidzos, President of RSA DSI, will present the Organizational Agreement.

Wednesday, March 13

9:30am IPSO/CIPSO BOF (Steve Crocker/TIS)

The status of the revised IP Security Option and the output of the Trusted Systems Interoperability Group (TSIG) Commercial IP Security Option will be discussed, as well as the documents themselves.

1:30pm Telnet (Dave Borman/Cray Research)

Members of the SAAG will be present at this meeting to discuss security issues.

7:00pm SNMP Security (James Galvin/TIS)

This will hopefully be the last meeting of this working group. The final version of the documents will be presented for approval. At the close of the meeting the documents will be submitted to the IESG for consideration as a proposed standard.

Thursday, March 14

9:00am Security Area Advisory Group meeting (Steve Crocker/TIS)

SAAG members are required to attend; all IETF attendees are welcome to observe. The status of the security related activities of the week will be summarized and plans for the next meeting will be discussed.

The Security Area within the IETF is responsible for development of security oriented protocols, security review of RFCs, development of candidate policies, and review of operational security on the Internet.

This report has two parts. The first section covers highlights from the meeting. The second section covers the organization and operation of the Security Area.

HIGHLIGHTS

Security Policy and Site Security Policy Handbook (SPWG and SSPHWG)

Both the Security Policy and Site Security Policy Handbook Working Groups prepared drafts of their documents. The security policy document is a concise statement of principles for protection of information assets and computing resources in the Internet. Because it's intended to act as a guide to others who will establish policies for their networks, hosts, products, etc., the IAB determined that this document will be called a Guidelines and will be issued as an informational RFC. The document is now available as an Internet Draft.

The Site Security Policy Handbook is an extensive document that is intended to serve as a basis for tailoring site-specific policies. It covers numerous facets of security including configuration, operation and responses to incidents.

These efforts are the result of the hard work and persistence of the Security Policy and Site Security Policy Handbook Working Groups. The members and particularly the Chairs of these groups deserve congratulations for the work they have done.

Common Authentication Technology (CAT)

John Linn and Jeff Schiller will co-Chair a new Working Group to explore and define a common authentication framework. This work will embrace MIT's Kerberos and Digital's SPx authentication servers. Digital also unveiled its General Security Services Application Program Interface (GSSAPI) which provides a common interface for SPx, Kerberos and any other authentication service that may be defined in the future. This work is intended to provide a uniform method for applications to authenticate connections in client-server and peer-peer connections.

Privacy Enhanced Mail (PEM)

The Privacy and Security Research Group (PSRG) under the Internet Research Task Force (IRTF) has revised the specifications for privacy enhanced mail. The specifications are being released as Internet Drafts and will be reviewed through the usual open process. At this IETF meeting, Jim Bidzos, the President of RSA Data Securityi, Inc, presented the outline of the forthcoming organizational agreement. (RSADSI holds the patent on the RSA public key technology and is licensing its use for privacy enhanced mail within the Internet.) Additional open meetings will be scheduled in forthcoming IETF meetings.

IP Security Option (IPSO)

Some time ago a protocol was defined for adding U.S. DoD security labels at the IP level. The protocol was never fully completed and sat in an incomplete state. Last fall, the effort was resurrected by Vint Cerf, the IAB Chair. Steve Kent has now completed the revisions to the document, and it is now available as an Internet Draft. This document covers only the Basic Security Option and is applicable only to the U.S. DoD security labels. Another document is expected later which will cover the Extended Security Option, and a separate effort is described next which is intended to cover labels outside of the U.S. DoD hierarchy.

Trusted Systems Interoperability Group (TSIG – CIPSO and TNFS)

The Trusted Systems Interoperability Group is a consortium of computer systems vendors developing protocols for trusted systems. Has asked the IETF and IAB for assistance in standardizing their protocols. The operation and rules of the TSIG are quite similar to the IAB and IETF. Each of the TSIG's protocols is developed by a TSIG Working Group whose deliberations are open to all. In order to facilitate the publication of protocols developed by the TSIG, the individual TSIG Working Groups will be chartered as IETF Working Groups. Two groups have submitted charters, CIPSO and TNFS.

The CIPSO Working Group is developing a commercial IP security option. This is intended to make security labels available to the commercial, civilian U.S. government and non-U.S. government communities. A draft document is essentially complete and will be made available as an Internet Draft.

The TNFS Working Group is developing a trusted version of the NFS (Network File System) protocol. This work is being coordinated with the distributed file systems Working Group in the Applications area. This work also depends on clarification of the status of NFS as a base for building other protocols.

ORGANIZATION AND OPERATION

Much of the work of the Security Area is performed in coordination with Working Groups in other areas. Indeed, one of the primary tasks is to provide security expertise to Working Groups in other IETF areas.

Starting with the December 1990 IETF meeting, we organized a Security Area Advisory Group (SAAG) to gather together the limited number of people knowledgeable about security in protocols and to provide a coordinated forum for discussion of security issues in Internet protocols. We've also established a pattern of having the SAAG meet twice during the IETF meeting, once at the beginning and once at the end of week. Although these are business meetings devoted principally to assignment of tasks and coordination of new work items, observers are welcome.

SAAG Operation

The main bulk of work for the SAAG consists of a set of formal work items. These work items correspond to three types of activities.

Security relevant developments within Working Groups in areas other than security.

Assistance to the Telnet Working Group on authentication and encryption is a typical example. For items of this type, a SAAG member is assigned and supports the Working Groups.

Working groups within the Security Area.

The development of SNMP security is an example. In many cases, even though a Working Group is in the Security Area, there are close ties to another area. SNMP security is obviously tied closely to the Management area. In several instances, it's a matter of choice whether a Working Group is in the Security Area or in another area. These decisions are made on a case by case basis by mutual agreement of the respective Area Directors. In these cases the work is usally coordinated closely with the relevant Area Director.

Preliminary inquiries

These are topics which do not merit the creation of a formal Working Group but which do need some level of attention. These are assigned to a SAAG member and followed for one or SAAG meeting.

In addition to the items formally being worked on by the SAAG, there are other discussions that take place but do not lead to the creation of a formal work item. No follow up actions are scheduled for these.

The following table shows the work items and other discussions arranged by status (SAAG, Security Area, Other Area, Prelim) and by which area they interact with. Minutes of the meetings of many of these groups are included in these proceedings.

	SAAG	Security Area	Other Areas	Prelim
Security	export iabcc	spwg		
Management		snmpsec		
User Services		ssphwg		
Routing			rreq	
Applications	passwd privdb chronos	cat pem(2)	telnet npp tnfs(1)	email nntp
Internet Services		ipso cipso(1)		iplpdn
OSI			ds	
Operations				

(1) This is a TSIG WG
(2) PEM is being developed by the PSRG

Trusted Systems
Interoperability Group

Jeffrey A. Edelheit

March 1991

What is the Trusted Systems Interoperability Group (TSIG)

- Consortium of vendors, systems integrators and government representatives
- Stated purpose is to develop interoperability specifications for trusted systems
- Primarily targeted towards UNIX servers and workstations
- Anyone may participate
- Meetings last 2.5 days
- Next meeting 7 - 9 May in Seattle

Who are TSIG Members

- Vendors
  - Digital Equipment
  - Sun Microsystems
  - Cray
  - Integraph
  - Silicon Graphics
  - SecureWare
  - Hewlett Packard
  - Harris
  - IBM
  - Wollongong Group
- Integrators
  - TRW
  - Loral Aerospace
- Government Representatives
  - MITRE
  - Sandia Nat'l Labs
  - Lawerence Livermore

TSIG Working Groups

- Commercial/Civil IP Security Option
- Trusted NFS
- Trusted Sockets
- Trusted X Windows
- Administration of Distributed Trusted Systems

Specifications Developed

- CIPSO
  - Requires Domain of Interpretation (DOI) Registrar
- Trusted NFS
  - Based on Sun's RPC, XDR protocols
  - Additional work needed on export and mount protocols

IETF Interests

- Establishing two IETF working groups now
  - CIPSO
  - TNFS
- Establish more IETF working groups when appropriate
- Move TSIG specifications into IAB standards process

E-Mail Addresses & Such

- Electronic Mailing lists
  - Adding to the primary list: tsig-request@wdl1.wdl.loral.com
  - General discussions: tsig@wdl1.wdl.loral.com
  - Subsidiary lists for CIPSO and TNFS
    - ❑ cipso or cipso-request @wdl1.wdl.loral.com
    - ❑ snfs or snfs-request@wdl1.wdl.loral.com
- Electronic Archives
  - archive-server @wdl1.wdl.loral.com
    "send index" on subject line

CURRENT MEETING REPORT

Reported by Jeff Schiller/MIT

PSWDMGMT BOF Minutes

Agenda

The Password Management group met in St. Louis as an informal Birds of a Feather session. The efforts of this group fall into two categories. The first is to define a series of protocols to be used within the context of the Common Authentication Technology (CAT). The second is to discuss issues involved in the use of one-time passwords on the Internet without the modification of existing TELNET client software.

It was decided that the first activity is best performed under the auspices of the CAT Working Group. The second activity isn't really an activity that would lead to a protocol per-se and is therefore best left to implementors. The SAAG will encourage such implementations and may in the future act to make the technology more widely available.

The group will cease formal activities. However, the exploration of one-time password mechanisms and their impact on telnet in particular will remain a SAAG work item.

Attendees

Karl Auerbach	karl@eng.sun.com
Richard Basch	probe@mit.edu
Warren Benson	wbenson@zeus.unomaha.edu
David Borman	dab@cray.com
Vinton Cerf	vcerf@nri.reston.va.us
Martina Chan	mchan@mot.com
Jeffrey Edelheit	edelheit@smiley.mitre.org
Barbara Fraser	byf@cert.sei.cmu.edu
Neil Haller	nmh@bellcore.com
J. Paul Holbrook	holbrook@cic.net
Joel Jacobs	jdj@mitre.org
Philip Karn	karn@thumper.bellcore.com
Steven Lendt	network@zeus.unomaha.edu
John Linn	linn@zendia.enet.dec.com
Bill Manning	bmanning@houston.sc.ti.com
Donald Merritt	don@brl.mil
Robert Reschly	reschly@brl.mil
Jeffrey Schiller	jis@mit.edu
Tim Seaver	tas@mcnc.org
Paul Tsuchiya	tsuchiya@bellcore.com

CURRENT MEETING REPORT

Reported by James Galvin/TIS

PEM I BOF Minutes

RFCs 1113, 1114 and 1115 are currently being revised. During this meeting, two new requirements that will be specified in the revised RFCs were discussed.

- Construction of distinguished names
- Implementation requirements

The revised RFCs, more precisely RFC 1114, will be specifying the use of mandatory, optional and prohibited attributes to be used in the construction of distinguished names. The reasons both for and against the rules were discussed. The principal reason in favor of the explicit rules was enhancement of the assurance of the proper operation of the certificate infrastructure. The ability to recognize and distinguish individual certificates from the those of a certification authority is a desirable feature. The principal reason against the explicit rules was the concern about interoperability with directory services pilots and the migration of PEM to use these services. However, the discussion did not yield a technical basis for the concern.

The revised RFCs, more precisely RFC 1114, will be specifying requirements on implementations that directly affect the user interface. Although it was agreed that security is enhanced if the requirements are implemented, it was unanimously agreed there are other mechanisms by which the "concepts" could be met. There was a good deal of concern about this issue. The Chair was tasked with bringing the issue to the attention of the Security Area Director and forwarding comments as appropriate.

Attendees

James Galvin	galvin@tis.com
Anthony Lauck	lauck@tl.enet.dec.com
John Linn	linn@zendia.enet.dec.com
E. Paul Love	loveep@sdsc.edu
Michael Reilly	reilly@pa.dec.com
Jeffrey Schiller	jis@mit.edu
Sam Sjogren	sjogren@tgv.com

CURRENT MEETING REPORT

Reported by James Galvin/TIS

PEM II BOF Minutes

D. James Bidzos, President of RSA DSI, presented a summary of the Organizational Agreement that will be required when PEM is made openly available to the Internet community. The status of the Organizational Agreement with respect to the IAB standards process was discussed, but not resolved. It was pointed out that previous versions of the agreement included security requirements. These requirements have been separated from the agreement and will be presented in a fifth PEM related RFC that will be published with the revised RFCs. The slides from Jim's presentation are included below.

Attendees

Jim Bidzos	
Vinton Cerf	vcerf@nri.reston.va.us
Stephen Crocker	crocker@tis.com
James Galvin	galvin@tis.com
Tom Kessler	kessler@sun.com
John Linn	linn@zendia.enet.dec.com
Carl Malamud	carl@malamud.com
Jeffrey Schiller	jis@mit.edu

TLCA RELATIONSHIPS

TLCA - ORGANIZATION AS ISSUER

- TLCA CERTIFIES ORG KEY
- ORG EMPLOYS CGU TO ISSUE CERTIFICATES
- ORG ISSUES CERTIFICATES

TLCA - ORGANIZATION CO-ISSUER

- TLCA OPERATES CGU ON BEHALF OF ORG
- ORGANIZATIONAL NOTARY (ON) REQUESTS CGU SERVICES VIA PEM
- CERTIFICATES RETURNED TO ON BY TLCA VIA PEM
- ORG APPEARS AS ISSUER IN THESE CERTIFICATES

TLCA-ORG AGREEMENT: PURPOSE

- VEHICLE BY WHICH AN ORG REQUESTS LICENSE TO ISSUE CERTIFICATES
- "STANDARD" TERMS UNDER WHICH SUCH LICENSE IS ISSUED

- DEFINE SERVICES PROVIDED BY TLCA
- PROVIDE LICENSE TO ORG TO:
  - ISSUE CERTIFICATES IN ACCORDANCE WITH RFC'S (ISSUER LICENSE)
  - PRACTICE PATENTED PUBLIC-KEY TECHNIQUES (PATENT LICENSE) IN CONJUNCTION WITH ISSUER LICENSE AND PRIVACY ENHANCED MAIL IN ACCORDANCE WITH RFC'S

TLCA RESPONSIBILITIES

- INVESTIGATION AND QUALIFICATION OF ORG AS ENTITY
- GRANTING OF PATENT/INTELLECTUAL PROPERTY RIGHTS
- NON-DISCRIMINATORY LICENSING AND LICENSE RENEWAL
- CERTIFICATION AND DELIVERY OF ORG KEY
- RESPONSIBLE MANAGEMENT OF TLCA KEY
- RESPONSIBLE MANAGEMENT OF CGU'S AS CO-ISSUER
- TIMELY AND ACCURATE ORG CRL REPORTING

ORGANIZATION RESPONSIBILITIES

- REPRESENTATION OF FACTUAL DATA
- DESIGNATION OF ON
- PAYMENT FOR CERTIFICATES ISSUED
- REVOKE CERTIFICATES WHERE CRITICAL INFORMATION HAS CHANGED
  - COMPLIANCE WITH RFC'S
  - RESPONSIBLE PRACTICES IN ISSUE OF CERTIFICATES
  - TIMELY AND ACCURATE CRL REPORTING
  - CERTIFY AFFILIATED USERS AND OU'S ONLY

RESIDENTIAL OR "NOTARY" USER CERTIFICATES

- WILL BE AVAILABLE FROM ORG OR ORGS LICENSED TO ISSUE NON-AFFILIATED CERTIFICATES
- USER SUBMITS NOTARIZED CERTIFICATE REQUEST ON PAPER AND KEY VIA EMAIL
- USER OBTAINS INDIVIDUAL PATENT LICENSE TO PRACTICE PEM IN CONJUNCTION WITH USE OF CERTIFICATE
- TERMS AND CONDITIONS SAME AS ORG AGREEMENT

PERSONA CERTIFICATES

- "ANONYMOUS" USERS
- ISSUE GUIDELINES (TBD) SIMILAR TO DMV PRACTICES

MISCELLANEOUS PEM/RSA/RSADSI ISSUES

- SCOPE OF PATENT RIGHTS
  - LIMITED TO USE IN CONJUNCTION WITH CERTIFICATE
  - CURRENTLY LIMITED FOR USE WITH NON-COMMERCIAL "INTERNET" EMAIL AND OTHER INTERNET PROTOCOLS AS MAY BE SPECIFIED IN FUTURE RFC'S
- NO RESTRICTIONS ON USE OF CERTIFICATE ITSELF - CAN BE USED WITH ANY LICENSED PRODUCT OBTAINED FROM ANY LICENSED VENDOR
- CERTIFICATE VALIDITY LIMITED TO TWO YEARS
- MD2 LICENSE IS IDENTICAL IN SCOPE TO RSA PATENT LICENSE
- MD4 HAS NO RESTRICTIONS OTHER THAN COPYRIGHT NOTICE REQUIREMENT AND MAY BE USED COMMERCIALLY

CURRENT MEETING REPORT

Reported by John Linn/DEC

CATBOF Minutes

A Birds of a Feather session met on Common Authentication Technology (CAT) at the March meeting; the first formal CAT Working Group meeting will take place at the July IETF. At the March BOF, Jeff Schiller and John Linn presented material on CAT concepts and responded to questions from attendees.

CAT's goal is to provide security services to a range of IETF protocol callers in a manner which insulates those callers from the specifics of underlying cryptographic security mechanisms, enabling modular separation between protocol and security implementation activities. Agreement on common security service interface characteristics, token representations, and other protocol integration issues, as well as discussion of individual mechanisms, falls within this Working Group's Charter. Two IETF applications protocol Working Groups (Telnet and Network Printing) are currently seeking to employ CAT-related techniques.

There was some controversy about mechanism type negotiation as contemplated by the Telnet security proposals. One observation: It's necessary to intersect two peers' notions of acceptable mechanisms, not for a client to accept any (however weak) which may be offered by a server. A belief was voiced that few servers would support more than a single mechanism, and/but that clients would often have to support multiple mechanisms to conform with their desired set of target servers; cases of single-mechanism clients communicating with multi-mechanism servers are also possible. While it was widely agreed that the world would be a better and more interoperable place if and when only one mechanism was in general use, there was a sense that ambidextrous hosts were unavoidable and would have to be accommodated. The Assigned Numbers RFC was proposed as a "registry" vehicle for mechanism type specifiers to be used in the Internet.

Interest was expressed in means to allow protection of data carried in stream-oriented protocols as well as in message-oriented protocols, whether by definition of stream-oriented security services interfaces or by (direct or mediated) provision of session keys to callers. There was debate about the merits of modeling protected password exchanges as CAT authentication mechanisms. In subsequent Security Area Advisory Group (SAAG) discussion, it was agreed that mechanisms performing key exchange, and hence constituting a basis for confidentiality and integrity protection for messages as well as authentication, should be emphasized.

The CAT activity will be supported with a family of documents, to be provided from different sources. A high-level Generic Security Service Application Program Interface (GSSAPI) specification will be submitted to the Internet-Draft process in advance of the July IETF meeting, and will be followed by a separate document defining a set of C language bindings therefore. Organizations defining particular security mechanisms (e.g., SPX, Kerberos) will submit separate mechanism-specific documents, supporting independently developed yet interoperable implementations of those mechanisms. CAT participants will pursue design refinements, protocol integration, and implementation activities, and will continue consulting liaison activities with IETF protocol Working Groups which are prospective clients for CAT-provided security services.

Attendees

Warren Benson	wbenson@zeus.unomaha.edu
Randy Butler	rbutler@ncsa.uiuc.edu
Vinton Cerf	vcerf@nri.reston.va.us
Martina Chan	mchan@mot.com
Stephen Crocker	crocker@tis.com
Jeffrey Edelheit	edelheit@smiley.mitre.org
Barbara Fraser	byf@cert.sei.cmu.edu
Shawn Gallagher	gallagher@quiver.enet.dec.com
James Galvin	galvin@tis.com
Tom Grant	grant@xylogics.com
Neil Haller	nmh@bellcore.com
Russ Hobby	rdhobby@ucdavis.edu
Joel Jacobs	jdj@mitre.org
Ajay Kachrani	kachrani@regent.enet.dec.com
Philip Karn	karn@thumper.bellcore.com
John Linn	linn@zendia.enet.dec.com
Mike Little	little@ctt.bellcore.com
Stephanie Price	price@cmc.com
Michael Reilly	reilly@pa.dec.com
George Sanderson	sanderson@mdc.com
Tim Seaver	tas@mcnc.org
Sam Sjogren	sjogren@tgv.com
Michael St. Johns	stjohns@umd5.umd.edu
William Townsend	townsend@xylogics.com
Glenn Trewitt	trewitt@pa.dec.com
Daniel Weidman	weidman@wudos2.wustl.edu

3.7.1 Internet Security Policy (spwg)

Charter

Chair(s):
Richard Pethia, `rdp@cert.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 (SPWG) 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, database 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.
Done	Revise and approve framework documents. Begin work on detailing areas of concern, technical issues, legal issues, and recommendations for each level of the hierarchy.
Done	Prepare first draft policy recommendation for Working Group review andmodification.
Sep 1990	Finalize draft policy and initiate review following standard RFC procedure.

CURRENT MEETING REPORT

Reported by Richard Pethia/CERT

SPWG Minutes

The Security Policy Working Group (spwg) met during the Twentieth Internet Engineering Task Force (IETF) meeting in St. Louis, on Tuesday, March 12, 1991. The latest draft of the Proposed Security Policy was presented and discussed.

Discussion during the meeting focused on two areas of concern: user authentication and local security.

User Authentication

While there is general agreement that individual users should be held accountable for their actions, there is not the same level of agreement that all users should be unambiguously identified for all types of Internet access.

Proponents of strong, mandatory, user authentication and access control mechanisms point to problems caused by "general use" accounts and "open" (without password) terminal servers where individuals take advantage of these open systems and use them as platforms to attack (access without authorization) other Internet systems. This group believes the use of simple user authentication and access control mechanisms would significantly reduce the problem. Steve Wolff, National Science Foundation (NSF), supported this position and indicated that it is NSF's position that individual user authentication and accountability should be required for access to NSFNET.

Opponents to this view believe enforced, unambiguous identification for all Internet access would potentially:

- Restrict the utility of the network (e.g., not allow a university library to set up "open" terminals that allow the university's students to browse the information resource),
- Place an administrative burden (e.g., issuing all university students unique account names and passwords, and managing those accounts and passwords) on sites that would be too expensive for some sites to bear,
- Infringe on a person's privacy by collecting data on the person's actions.

Rather than attempt to resolve the controversy at this point in time, it was decided that the proposal would be changed to remove the phrases that called for a ban on "open" servers and stress the importance of individuals' accountability for their actions.

Local Security

Another area of concern was the elaboration section of item 3 (local security). Included in this section was a listing of five elements needed for good local security. This listing treated local security in greater depth than any other issue in the document. To balance the discussion of issues, the list was removed from the body of the proposal and included as an appendix. In addition, it was decided that two of the elements listed would be modified according to suggestions and comments received. The group discussed that there are trade-offs between strict security and the usability of systems. A paragraph would be added to touch on this subject.

Additional discussion centered around how the document would be used and interpreted. Some people felt that since the title included the word "policy", it would be used as if it were legally enforceable. For this reason the title of the document was changed to "Guidelines for the Secure Operation of the Internet". Necessary changes within the body of the document would be made to match the title change.

The group felt that it was necessary to push forward with the document. Vint Cerf suggested that the nature of this document was unique within the document collection of the IETF and that it would be helpful to have it reviewed by the Internet Advisory Board (IAB). The IAB could then advise the group as to how the document should be handled. To that end, the following schedule was set.

March 18	Final draft completed
March 19	Draft emailed to internet-drafts@nri.reston.va.us
April 3	Document to be discussed during IAB teleconference.

Whether or not the Working Group meets at the next IETF will be based upon the outcome of the IAB's review of the document.

Attendees

Warren Benson	wbenson@zeus.unomaha.edu
David Benton	benton@bio.nlm.nih.gov
Randy Butler	rbutler@ncsa.uiuc.edu
Vinton Cerf	vcerf@NRI.Reston.VA.US
Martina Chan	mchan@mot.com
Stephen Crocker	crocker@tis.com
Jeffrey Edelheit	edelheit@smiley.mitre.org
Fred Engel	engel@concord.com
Barbara Fraser	byf@cert.sei.cmu.edu
Neil Haller	nmh@bellcore.com
Sergio Heker	heker@jvnc.net
J. Paul Holbrook	holbrook@cic.net
Philip Karn	karn@thumper.bellcore.com
April Merrill	
Richard Pethia	rdp@cert.sei.cmu.edu

Robert Reschly	reschly@brl.mil
Jeffrey Schiller	jis@mit.edu
Tim Seaver	tas@mcnc.org
Albert Soule	als@sei.cmu.edu
Mike Turico	mturico@mot.com
Daniel Weidman	weidman@wudos2.wustl.edu
Stephen Wolff	steve@nsf.gov
C. Philip Wood	cpw@lanl.gov
Osmund deSouza	desouza@osdpc.ho.att.com

3.8 User Services Area

Director(s):

- Joyce Reynolds: jkrey@isi.edu

Area Summary reported by Joyce Reyolds/ISI

Three new Working Groups met at the IETF in St. Louis:

- Directory Information Services (pilot) Infrastructure Working Group (DISI), Chaired by Chris Weider.

- Internet User Glossary Working Group (userglos) Chaired by Karen Roubicek and Tracy LaQuey Parker

- NOC-Tool Catalogue Revisions Working Group (noctool2) Chaired by Robert Enger and Gary Malkin

Working Group Reports

DISI - Chaired by Chris Weider

DISI is chartered to facilitate the deployment in the Internet of Directory Services based on implementations of the X.500 standards. It will facilitate this deployment by producing informational RFCs intended to serve as a Directory Services "Administrator's Guide". These RFCs will relate the current usage and scope of the X.500 standard and Directory Services in North America and the world, and will contain information on the procurement, installation, and operation of various implementations of the X.500 standard. As the various implementations of the X.500 standard work equally well over TCP/IP and CLNP, the DISI Working Group shall not mandate specific implementations or transport protocols.

DISI is an offshoot of the OSI Directory Services group and is a combined effort of the User Services Area and the OSI Integration Area of the IETF.

NISI - Chaired by Dana Sitzler and Patricia Smith

The "Building a Network Information Services Infrastructure" draft document was expanded to more accurately define and describe a NIC, with the ultimate goal intended to make it easier for users to get information from NICs. The current Working Group draft was gone over with the NISI members, with the intent to install it as an Internet Draft after the St. Louis IETF.

NOCTOOL2 - Chaired by Robert Enger and Gary Malkin

A new Charter is in place. The "Son of NOCTools" Working Group is currently in the process of updating and revising their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

USERGLOS - Chaired by Karen Roubicek and Tracy LaQuey Parker

The User Gloss Working Group is chartered to create an Internet glossary of networking terms and acronyms for the Internet community. At this meeting the Working Group went over and amended the Charter, defined the criteria for glossary terms, reviewed existing glossaries, discussed glossary format, and discussed and defined processes for drafting, editing and group review.

SSPHWG - Chaired by J. Paul Holbrook and Joyce K. Reynolds

Please consult the Security Area Director's report for information on this group's current progress. The SSPHWG is a combined effort of the Security and User Services Area of the IETF.

USWG - Chaired by Joyce K. Reynolds

Agenda items included:

- QUAIL - presented by Gary Malkin

  Gary led a brief discussion of the two currently issued Quail documents, and requested additional volunteers to continue to help monitor the mailing lists, and to contribute to future updates of Quail.

- DISI - presented by Chris Weider
  (See above working group report.)

- The USWG published two documents:

  - Malkin, G., and A. Marine, "FYI on Questions and Answers - Answers to Commonly asked 'New Internet User' Questions", FYI 4, RFC 1206.

  - Malkin, G., Marine A., and J. Reynolds, "FYI on Questions and Answers - Answers to Commonly asked 'Experienced Internet User' Questions", FYI 7, RFC 1207.

Announcement

The User Services Area of the Internet Engineering Steering Group (IESG), has established a User Services Area Council (USAC) to promote and encourage creative exchange of international user service needs and concepts. Constructive input from various national and international user services organizations for the purpose of not duplicating each organization's efforts is also encouraged. USAC will be responsible for researching and defining short-term and long-term user services needs and coordinating developments in finding solutions.

The primary responsibilities of USAC members are to actively provide input for the current and future developments of user services concerns. This forum will conduct meetings in conjunction with the IETF plenaries. Primary interaction among members will take place via electronic mail. The User Services Area Director of the IESG chairs the USAC. Current membership includes representation from Australia, Canada, Europe, Israel and Japan.

3.8.1 Directory Information Services Infrastructure (disi)

Charter

Chair(s):
Chris Weider, clw@merit.edu

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

Description of Working Group:

The Directory Information Services (pilot) Infrastructure Working Group (DISI) is chartered to facilitate the deployment in the Internet of Directory Services based on implementations of the X.500 standards. It will facilitate this deployment by producing informational RFCs intended to serve as a Directory Services "Administrator's Guide". These RFCs will relate the current usage and scope of the X.500 standard and Directory Services in North America and the world, and will contain information on the procurement, installation, and operation of various implementations of the X.500 standard. As the various implementations of the X.500 standard work equally well over TCP/IP and CLNP, the DISI Working Group shall not mandate specific implementations or transport protocols.

The Directory Information Services (pilot) Infrastructure Working Group is an offshoot of the OSI Directory Services group, and, accordingly, is a combined effort of the OSI Integration Area and User Services Area of the IETF. The current OSIDS Working Group was chartered to smooth out technical differences in information storage schema and difficulties in the interoperability and coherence of various X.500 implementations. The DISI group is concerned solely with expanding the Directory Services infrastructure. As DISI will be providing infrastructure with an eye towards truly operational status, DISI will need to form liasons with COSINE, Paradyse, and perhaps the RARE WG3.

As a final document, the DISI Working Group shall write a Charter for a new Working Group concerned with user services, integration, maintenance, and operations of Directory Services, the Internet Directory User Services Group.

Goals and Milestones:

Mar 1991 First IETF Meeting: review and approve the Charter making any changes necessary. Examine needs and resources for the documentation to be produced, using as a first draft a document produced by Chris Weider, MERIT, which will be brought to the IETF. Assign writing assignments. Further work will be done electronically.

Jul 1991 Second IETF Meeting: review and approve documentation; review and approve Charter for the IDUS group.

Aug 1991 Electronically review final draft of documentation, and, if acceptable, submit to IESG for publication.

Dec 1991 Third IETF Meeting: Declare success and reform DISI group as IDUS group.

CURRENT MEETING REPORT

Reported by Joyce Reynolds/ISI

DISI Minutes

The Distributed Information Service Infrastructure Working Group met during the User Services Working Group. The Chair, Chris Weider led the session which included discussion on the current Charter, current deployment and players, available implementations, documents to be produced, and how this group segues into User Services activities.

In reviewing the current Charter, suggestions were made to change the last paragraph to include "maintenance", and in the 2nd paragraph, change "building" to "providing" information.

Chris will amend the Charter, and send out a new one for approval. The "combined efforts" of the OSI Integration Area and the User Services Area was discussed, as was what the current deployment of DSAs are. There are currently more than 200 DSAs in 20 countries!

What organizations are currently building infrastructure? The FOX project at ISI, Merit, PSI, and SRI. In Europe, PARADISE. Quipu - Part of the ISODE, IAN at UBC, NIST (not fully functional yet).

Documents to be produced by DISI:

- Vol. 1 - Advantage Document; What should I do with it, even if I know what it is??
- Vol. 2 - Implementation Document; rapidly changing. Would provide current implementation and interoperability profiles.
- Vol. 3 - Advanced usages document/administrators guide (1 year life span).

Writing Assignments:

- Vol. 1 - Chris Weider, Sergio Heker and Joyce Reynolds
- Vol. 2 - Ruth Lang and Russ Wright
- Vol. 3 - Assignments will be made at the next IETF in Atlanta.

Attendees

Richard Bowles	bowles@stsci.edu
Robert Enger	enger@seka.scc.com
Charles Fumuso	cwf@cray.com
Keith Hacke	hacke@informatics.wustl.edu
Sergio Heker	heker@jvnc.net
Darren Kinley	kinley@crim.ca
Tracy LaQuey Parker	tracy@utexas.edu
Ruth Lang	rlang@nisc.sri.com
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@bbn.com
Dana Sitzler	dds@merit.edu
Patricia Smith	psmith@merit.edu
Michael Stanton	"usermas@lncc.bitnet"
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Joanie Thompson	joanie@nsipo.nasa.gov
Chris Weider	clw@merit.edu
Russ Wright	wright@lbl.gov

3.8.2 Internet User Glossary (userglos)

Charter

Chair(s):

Karen Roubicek, roubicek@bbn.com
Tracy Parker, tracy@utexas.edu

Mailing Lists:

General Discussion: usergloss@ftp.com
To Subscribe: usergloss-request@ftp.com

Description of Working Group:

The User-Gloss Working Group is chartered to create an Internet glossary of networking terms and acronyms for the Internet community.

Goals and Milestones:

Done	Examine the particular Internet user needs for a glossary and define the scope. Review, amend, and approve the Charter as necessary. Discussion of Userglos Working Group Chair nominations submitted by USWGers.
TBD	Review Internet user needs and format for a glossary. Discussion of current ideas about the glossary and the outline development. Finalize outline and organization of the glossary.
TBD	Draft of glossary will be prepared, draft to be reviewed and modified.
TBD	Second pass draft of glossary. Draft to be reviewed and modified, finalize draft glossary.
TBD	Initiate IETF Internet Draft review process by submission of Userglos draft to IETF Secretary. Follow-up with the submission of the glossary to RFC Editor as an FYI RFC.

CURRENT MEETING REPORT

Reported by Tracy LaQuey Parker/UTexas

USERGLOS Minutes

Status Update

Chairpersons	Karen Roubicek / roubicek@nnsc.nsf.net Tracy LaQuey Parker / tracy@utexas.edu
Mailing List	usergloss@ftp.com subscription - usergloss-request@ftp.com
Date of Last Meeting	St. Louis IETF / 11 March 1991
Date of Next Meeting	Atlanta IETF / August 1991
Pending/New Objectives	Compile list of terms, acronyms and definitions from existing sources. Will then edit this list.
Progress to Date	Have compiled a list of existing glossaries to use as a base.

Agenda

- Review Charter
- Define Criteria for Glossary Terms
- Review Existing Glossaries
- Discuss Glossary format
- Discuss online version
- Establish processes for drafting, editing, and group review
- Review proposed milestones and adjust
- List action items for summer IETF and next 4 months

This was the first official full meeting of the USERGLOSS Working Group. The meeting began with a review of the Charter. To summarize, the objective of this group is to create a glossary of networking terms and acronyms for the Internet community. There was some discussion about creating another glossary when several already exist, and one has been published as RFC 1208 (Glossary of Networking Terms). However, the general feeling was that there is no complete glossary that truly represents the Internet community.

The criteria for selection of glossary terms was discussed. Because we don't have anything to work with yet, we decided to postpone this step until the next meeting. There was some discussion on how to decide which organizations would be included as entries. The general consensus was to admit all organizations that directly support the Internet. Some criteria that was suggested included, federal agencies; companies who have contracts in support of government or networking (such as ISI, SRI); state, regional, and midlevel networks;

backbone networks; non-profit organizations; Internet specific organizations (IAB, IESG). We will not include names of universities or companies.

Other issues were brought up. Should we limit the glossary to the Internet community? It was decided that the glossary should reflect our perspective FROM the Internet community. The scope of the glossary will include IETF terms and terms used in required RFC's. The audience for this FYI will be the broad range of Internet users.

We reviewed a list of existing glossaries. After looking at this list Ole Jacobsen observed that "There are too many words here." The list included:

- Networking.terms found on nsipo.arc.nasa.gov in pub directory.
- The glossary in RFC 1206, "Answers to Commonly asked 'New Internet User' Questions".
- RFC 1208, "A Glossary of Networking Terms".
- Glossary in the NNSC Internet Tour.
- Glossary from "Analyzing Sun Networks" by Carl Malamud.
- Glossaries from other textbooks, such as Comer's TCP/IP and Rose' books.
- Hacker's Dictionary.
- NCAR's Glossary.

We then talked about the format of the glossary. It will be ascii text. For now, Tracy will keep the glossary in a standard format that can be converted to what we decide on later. A suggested format was:

> term/acronym/expansion — Definition (1,2,3...) — source of definition — index

Indexes will be defined later. These were suggested:

- Protocol
- Acronym
- Operating networks
- Organizations
- Government
- International
- Level (new user, etc.).

The online draft version will be kept on a host at the University of Texas. Tracy will announce where it is to the Usergloss mailing list. It was mentioned that we should think about generating several forms of the glossary - a short form and an expanded form.

Next on the agenda was the editing process. It was decided that the approach will be to create a "big bucket" of existing online glossaries (and include attributions when appropriate). An editing committee, consisting of Gary Malkin, Tracy LaQuey Parker, Karen Roubicek, and Carl Malamud, will edit this "bucket" and suggest additions and deletions. Two people that volunteered at the Boulder IETF who were not present at this meeting are Marilyn Martin and Allen Apt.

The following editing steps were suggested (these are also our proposed milestones):

Step 1 Compile from existing sources. Remove duplicates and definitions that don't apply. Complete by August 1991

Step 2 Submit for group review. Resolve any conflicts. Add new terms. Refine indexes. To be done at August IETF.

Step 3 Assign indexes/classify terms. August - September 1991

Step 4 Group Review October - November 1991

Step 5 Submit as Internet Draft December 1991

And finally, we made a list of action items:

- Gary Malkin will update and provide for archiving of the mailing list.
- Tracy LaQuey Parker will find a host machine for the glossary and announce it to the mailing list.
- Karen Roubicek will follow up on the two people on the editorial board who were not present at this meeting.
- Someone needs to research the "Hacker's Dictionary."
- Karen will contact Don Morris about the NCAR glossary.
- Gary Malkin will send the online glossary from RFC 1206 to Tracy.
- Accomplish Step 1 defined above (editorial board will divide up these tasks).

Attendees

Robert Enger	enger@seka.scc.com
Douglas Gale	dgale@note.nsf.gov
Sergio Heker	heker@jvnc.net
Ole Jacobsen	ole@csli.stanford.edu
Shelly Knueven	shelly@wugate.wustl.edu
Tracy LaQuey Parker	tracy@utexas.edu
Steven Lendt	network@zeus.unomaha.edu
Carl Malamud	carl@malamud.com
Gary Malkin	gmalkin@ftp.com

Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@bbn.com
Patricia Smith	psmith@merit.edu
Joanie Thompson	joanie@nsipo.nasa.gov
John Wobus	jmwobus@suvm.acs.syr.edu

3.8.3 NOC-Tool Catalogue Revisions (noctool2)

Charter

Chair(s):

Robert Enger, enger@seka.scc.com
Gary Malkin, gmalkin@ftp.com

Mailing Lists:

General Discussion: noctools@merit.edu
To Subscribe: noctools-request@merit.edu

Description of Working Group:

The NOC-Tools Working Group will update and revise their catalog to assist network managers in the selection and acquisition of diagnostic and analytic tools for TCP/IP Internets.

- Update and revise the reference document that lists what tools are available, what they do, and where they can be obtained.
- Identify additional tools available to assist network managers in debugging and maintaining their networks that were inadvertently omitted in previous NOCTools catalog.
- Identify additional new or improved tools that have become apparent since the last the compilation of the reference document.
- Arrange for the central (or multi-point) archiving of these tools in order to increase their availability.
- Establish procedures to ensure the ongoing maintenance of the reference and the archive, and identify an organization willing to do it.

Goals and Milestones:

Mar 1991 Review Internet tool needs and updates/corrections for the "Son of NOCTools" catalog. Discussion of additional input to the catalog.

Aug 1991 Draft of catalog will be prepared, draft to be reviewed and modified.Initiate IETF Internet Draft review process by submission of a "Son of NOCTools" catalog draft to IETF Secretary.

Dec 1991 Follow-up with final amendments to the document and the submission of the catalog to RFC Editor as an FYI RFC for publication

CURRENT MEETING REPORT

Reported by Gary Malkin/FTP Software and Bob Enger/Contel

NOCtools2 Minutes

Agenda

- Suggested format changes
- Suggested augmentations (indexes, charts, etc)
- Updating entries
- On-going maintenance
- Locating worker bees (appointing volunteers)

Discussions

- Format Changes:

  Everybody seemed happy with the current format so it was decided to leave it as it is.

- Augmentations:

  If there is always a postscript version, we should have a nice cross-reference chart for keywords and tools. This was discussed in NOCtools originally.

- Updating Existing Entries:

  It was decided that the contacts for the entries in the current catalog will be contacted for updates. In the future, all entries will have a "last updated" timestamp. Entries which are not updated by the contacts within some (to be determined) time, will be deleted. Updates must be submitted as whole replacements for entries. Requests to modify an entry will be refused (too much overhead).

- Adding New Entries:

  We will send mail to ietf, tcp-ip, pc-ip and snmp mailing lists announcing the re-opening of the catalog. Other announcement possibilities are ConneXions and trade magazines (which didn't work too well last time). Additionally, it was decided to poll the user community (rather than just vendors) to ask what tools they use often (NB: care must be taken to avoid words like "recommended").

- On-going Maintenance:

  This is a topic which plagues the User Services Working Group since many documents are "living documents". Various suggestions, such as on-line databases, were discussed; however, no consensus was reached.

Other Items

- The Charter needs to be tweeked to be NOCtool2.
- Contact John Wobus about big-lan mailing list.
- Contact Don Morris about graduate student tools.
- Cliff Newman suggested Prospero for remote access.
- Is ODA useful?
- Check in with Network Distributed Database Working Group.
- Jim Sheridan suggests dial-up modem to upload new entries.

Action Items

Bob Enger	Divide current catalog among himself, Gary Malkin and Darren Kinley to get tool updates.
Gary Malkin	Compose user community poll and send to NOCtools list for refinement.

Attendees

Steve Bostock	steveb@novell.com
Richard Bowles	bowles@stsci.edu
Fred Engel	engel@concord.com
Robert Enger	enger@seka.scc.com
Bill Fardy	fardy@ctron.com
Douglas Gale	dgale@note.nsf.gov
Shawn Gallagher	gallagher@quiver.enet.dec.com
Keith Hacke	hacke@informatics.wustl.edu
Mike Janson	mjanson@mot.com
Kenneth Key	key@cs.utk.gdy
Darren Kinley	kinley@crim.ca
Cheryl Krupczak	clefor@secola.columbia.ncr.com
Gary Malkin	gmalkin@ftp.com
Joyce Reynolds	jkrey@isi.edu
Ron Roberts	roberts@jessica.stanford.edu
Kary Robertson	kr@concord.com.kr
Karen Roubicek	roubicek@bbn.com
Jonathan Saperia	saperia@decwrl.enet.dec.com
Jim Sheridan	jsherida@ibm.com
Patricia Smith	psmith@merit.edu
Michael Stanton	"usermas@lncc.bitnet"
Shujiuan Wang	swang@ibm.com
Daniel Weidman	weidman@wudos2.wustl.edu
Wing Fai Wong	wfwong@malta.sbi.com

3.8.4 Network Information Services Infrastructure (nisi)

Charter

Chair(s):

Dana Sitzler, dds@merit.edu
Pat Smith, psmith@merit.edu

Mailing Lists:

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

Description of Working Group:

The NISI Working Group 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 Working Group 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

NISI Minutes

Agenda

- Review Draft Document
- Discuss NIC Forum
- Discuss Next Steps

The meeting began with a review of the draft document. The goal is to submit the document to the RFC editor prior to tne next IETF meeting in July. The following revisions were suggested and will be made prior to submission:

1. Acknowledge the need for the registration function – but do not include as a NIC requirement for all NICs. Treat as a special NIC function.

2. NIC accountability information should include:
   - Time stamp.
   - Revision number.
   - If NIC produced, should indicate in document.
   - Information source is optional in file – NIC should have contact information but it does not necessarily have to go in the file.

3. The document should include examples of the nic@domain naming convention. Examples should indicate what user can expect if:
   - There is one nic at the domain and
   - There are multiple nics at the domain.

4. The NIC forum will be open to all NIC personnel. NIC profile information collected will be kept online as a resource and also included in the Internet Resource Guide.

5. Document language should be revised to be more inclusive of international and peer network structures. The document currently assumes the hierarchical structure of NSFNET. An attempt will be made to make such references more generic so they apply to other structures.

6. The document should acknowledge the relationship between NICs and NOCs

The majority of the meeting was spent discussing these document revisions. We also got into a lively discussion about whether this document addressed the needs of a network information services infrastructure. It was acknowledged that this document is only a first step in the process – an attempt to start the processing of setting NIC conventions. We also discussed the upcoming NSFNET User Services solicitation.

The other Agenda items were discussed briefly. The NIC forum will be implemented at Merit and all NICs will be encouraged to participate. Time ran out before the group could discuss what the nisi group should do next. This topic will be discussed on the mailing list.

Attendees

Kathy Atnip	kathy@wugate.wustl.edu
Richard Bowles	bowles@stsci.edu
Sean Donelan	sean@dra.com
Elizabeth Feinler	
Douglas Gale	dgale@note.nsf.gov
Fred Gray	fred@homer.msfc.nasa.gov
Keith Hacke	hacke@informatics.wustl.edu
Martyne Hallgren	martyne@theory.tn.cornell.edu
Ajay Kachrani	kachrani@regent.enet.dec.com
Darren Kinley	kinley@crim.ca
Tracy LaQuey Parker	tracy@utexas.edu
Ruth Lang	rlang@nisc.sri.com
Gary Malkin	gmalkin@ftp.com
Robert Morgan	morgan@jessica.stanford.edu
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@bbn.com
Dana Sitzler	dds@merit.edu
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Joanie Thompson	joanie@nsipo.nasa.gov
Chris Weider	clw@merit.edu
Wengyik Yeong	yeongw@psi.com

3.8.5 Site Security Policy Handbook (ssphwg)

Charter

Chair(s):

J. Paul Holbrook, holbrook@cic.net
Joyce K. Reynolds, jkrey@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:

1. 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.
2. Establishing procedures to prevent security problems:
   - System security audits.
   - Account management procedures.
   - Password management procedures.
   - Configuration management procedures.
3. 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).
4. Establishing post-incident procedures
   - Removing vulnerabilities.
   - Capturing lessons learned.
   - Upgrading policies and procedures.

Goals and Milestones:

Done — Review, amend, and approve the Charter as necessary. Examine the particular customer needs for a handbook and define the scope. Continue work on an outline for the handbook. Set up an SSPHWG "editorial board for future writing assignments for the first draft of document.

Done — Finalize outline and organization of handbook. Partition out pieces to interested parties and SSPHWG editorial board members.

Done — Pull together a first draft handbook for Working Group review and modification.

Oct 1990 — Finalize draft handbook and initiate IETF Internet Draft review process, to follow with the submission of the handbook to the RFC Editor for publication.

Oct 1990 — Finalize draft handbook and initiate IETF Internet Draft review process, to follow with the submission of the handbook to the RFC Editor forpublication.

CURRENT MEETING REPORT

Reported by Joyce K. Reynolds/ISI and J. Paul Holbrook/CERT

SSPHWG Minutes

Discussion of Handbook's current draft status from November 29, 1990. It has been decided to go ahead with the I-D process to RFC publication.

We have formally requested a 30 minute plenary slot at the Atlanta IETF to talk about the Site Security Policy Handbook, which we expect to be a FYI/RFC by that plenary.

We have established the following firm dates for the Handbook:

- The document is open for comments and revisions until May 15, 1991. Paul Holbrook will be working on the contents of the document for fine tuning and clarity per comments received back from SSPHWGers. Paul's deadline is April 15, 1991.

  Joyce Reynolds will take over and make editorial adjustments to the document, per comments received back from SSPHWGers on how the document could have a more even "flow" if some sections were rearranged. Joyce's deadline is May 15, 1991.

- After the May 15th close of comments and revisions, Paul and Joyce will continue to incorporate any final comments. The document will be submitted as an Internet-Draft by EOB Pacific Time, May 31, 1991.

- There were a substantial number of comments about Section 2, especially from Rich Pethia of the CERT. Barbara Fraser of CERT has agreed to take over Chapter 2, with the assistance of Allen Sturtevant, to rework it to be consistent with the comments received. An updated Chapter 2 will be provided by the May 15th deadline.

- Once the Handbook is in the I-D process, there will be a 30 day stay in I-D land, allowing for final comments by June 28, 1991. Paul and Joyce will edit any comments and submit the I-D to the RFC Editor by EOB Pacific Time, July 12, 1991.

Additional items:

- Ask Noel Chiappa about his lawyer report to put in the Handbook's Bibliography - per Vint Cerf suggestion.

- How about an Index in this Handbook?? Yes, per Vint Cerf. Joyce will check on this in the ms macros.

- PostScript/NRoff help will be provided by Vint Cerf and Jeff Edelheit.

- "Beta testing" the Handbook:
  - Allen Sturtevant will provide input from Bob Aiken
  - CERTTools
  - Washington University - Dan Weidman
  - Security Policy and SSPHWG members
  - Mike Roberts - EDUCOM

- Suggestion from Vint Cerf - organize a workshop that would use this document. CNRI could provide some sponsorship.

Attendees

Warren Benson	wbenson@zeus.unomaha.edu
Vinton Cerf	vcerf@NRI.Reston.VA.US
Martina Chan	mchan@mot.com
Stephen Crocker	crocker@tis.com
Jeffrey Edelheit	edelheit@smiley.mitre.org
Barbara Fraser	byf@cert.sei.cmu.edu
Martyne Hallgren	martyne@theory.tn.cornell.edu
Sergio Heker	heker@jvnc.net
J. Paul Holbrook	holbrook@cic.net
Joel Jacobs	jdj@mitre.org
Steven Lendt	network@zeus.unomaha.edu
Bill Manning	bmanning@houston.sc.ti.com
Mike Marcinkevicz	mdm@calstate.edu
Richard Pethia	rdp@cert.sei.cmu.edu
Joyce Reynolds	jkrey@isi.edu
Tim Seaver	tas@mcnc.org
Albert Soule	als@sei.cmu.edu
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Joanie Thompson	joanie@nsipo.nasa.gov
Daniel Weidman	weidman@wudos2.wustl.edu

3.8.6 User Services (uswg)

__Charter__

Chair(s):

Joyce K. Reynolds, jkrey@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 Working Groups 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);
  - Not duplicate existing or previous efforts.
- Create Working Groups 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

USWG Minutes

Announcements

- New Working Group - Internet User Glossary (USER-GLOSS)
- New Working Group - "Son of NOCTOOLS" (NOCTOOL2)
- New Working Group - Directory Information Services Infrastructure (DISI)
- Q/A for New Internet Users - published (RFC 1206, FYI 4)
- Q/A for Experienced Internet Users - published (RFC 1207, FYI 7)
- SSPHWG - met Tuesday afternoon, 3/12, 1:30pm-3:30pm
- NISI - met Wednesday afternoon, 3/13, 1:30pm-3:30pm

Discussions/Reports

QUAIL - presented by: Gary Malkin
Gary led a brief discussion of the two currently issued Quail documents, and requested additional volunteers to continue to help monitor the mailing lists, and to contribute to future updates of Quail.

DISI - presented by: Chris Weider
The rest of the USWG session was devoted to the new Working Group, Directory Information Services Infrastructure (DISI). Chris Weider led the session. Additional information can be obtained under the DISI Working Group Minutes.

Attendees

Joe Blackmon	blackmon@ncsa.uiuc.edu
Richard Bowles	bowles@stsci.edu
Sean Donelan	sean@dra.com
Robert Enger	enger@seka.scc.com
Charles Fumuso	cwf@cray.com
Douglas Gale	dgale@note.nsf.gov
Keith Hacke	hacke@informatics.wustl.edu
Sergio Heker	heker@jvnc.net
J. Paul Holbrook	holbrook@cic.net
Darren Kinley	kinley@crim.ca
Tracy LaQuey Parker	tracy@utexas.edu
Ruth Lang	rlang@nisc.sri.com
Gary Malkin	gmalkin@ftp.com
Joyce Reynolds	jkrey@isi.edu
Karen Roubicek	roubicek@bbn.com

Dana Sitzler	dds@merit.edu
Patricia Smith	psmith@merit.edu
Michael Stanton	"usermas@lncc.bitnet"
Allen Sturtevant	sturtevant@ccc.nmfecc.gov
Joanie Thompson	joanie@nsipo.nasa.gov
Chris Weider	clw@merit.edu
Russ Wright	wright@lbl.gov

Chapter 4

Network Status Briefings

4.1 Research Networking in Brazil Report

Presented by Michael Stanton/PUC-Rio Departamento De Informatica

RESEARCH NETWORKING IN BRAZIL

IETF ST. LOUIS MARCH 14, 1991

TADAO TAKAHASHI (EXECUTIVE COORD.)
TADAO@ETHOS1.ANSP.BR

DEMI GETSCHKO (OPERATIONS COORD.)
DEMI@BRFAPESP.BITNET

MICHAEL STANTON (R & D COORD.)
USERMAS@LNCC.BITNET

REDE NACIONAL DE PESQUISA/CNPq

BRAZIL

SIZE: ~ CONTINENTAL U.S.A.
(N ⟷ S ~ E ⟷ W ~ 2000 MILES)

LANGUAGE: PORTUGUESE

CAPITAL: BRASILIA

POPULATION: ~ 140,000,000

MAIN CITIES:

SAO PAULO	12,000,000
RIO DE JANEIRO	8,000,000
BELO HORIZONTE	1,500,000
BRASILIA	1,500,000
RECIFE	1,500,000
PORTO ALEGRE	1,000,000

TELECOM:

NAT. PTT (EMBRATEL) + REGIONALS

- MICROWAVE
- SOME FIBER (SP, RJ)
- UP TO 9.6 KBPS
  (1992 → 64 KBPS)
- PUBLIC X.25 (SINCE 1982)
- SATELLITE

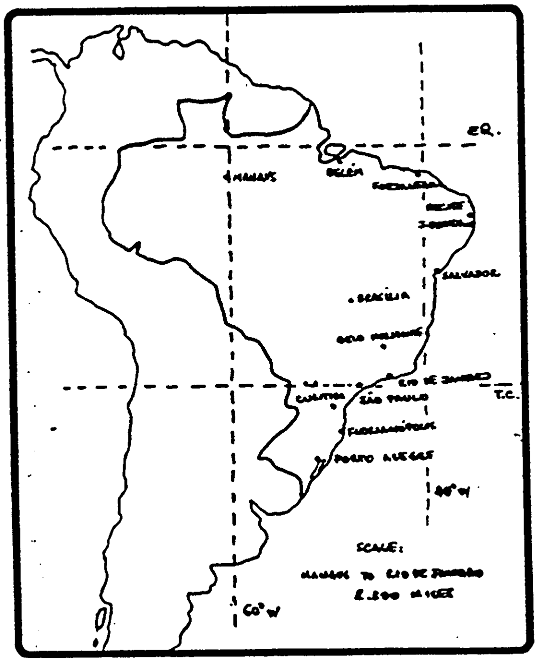

NETWORKING HISTORY

I. UNTIL 10/88 PTT DID NOT PERMIT SHARED USE OF (INTERNATIONAL) CIRCUITS

⟹ INDIVIDUAL INITIATIVES BY SEPARATE ORGANIZATIONS

09/88 LNCC (RIO) - UMD (BITNET, 9600)
11/88 FAPESP (SP) - FERMILAB (HEPNET + BITNET, 9600)
05/89 UFRJ (RIO) - UCLA (BITNET, 4800)

II. SHARED USE IS NOW "TOLERATED"

⟹ INTERCONNECTION OF AREAS SERVED BY DIFFERENT INTERNATIONAL CIRCUITS

04/90 LNCC - FAPESP
09/90 LNCC - UFRJ

BY 01/91 ABOUT 50 INSTITUTIONS INTERCONNECTED BY MIX OF LEASED CIRCUITS, X.25 PDN AND DIALUP.

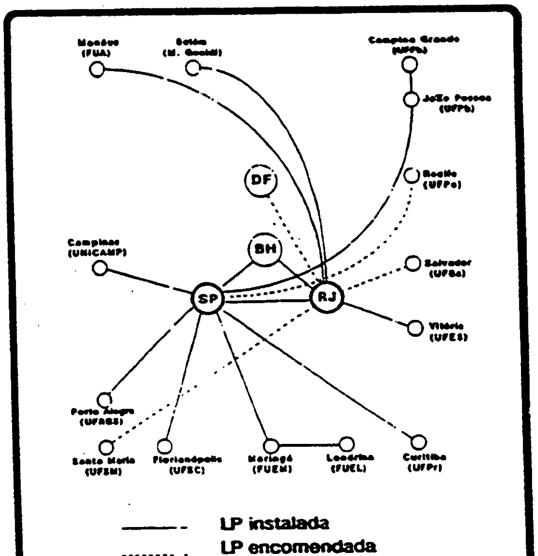

Conexões da RNP (jan/91)

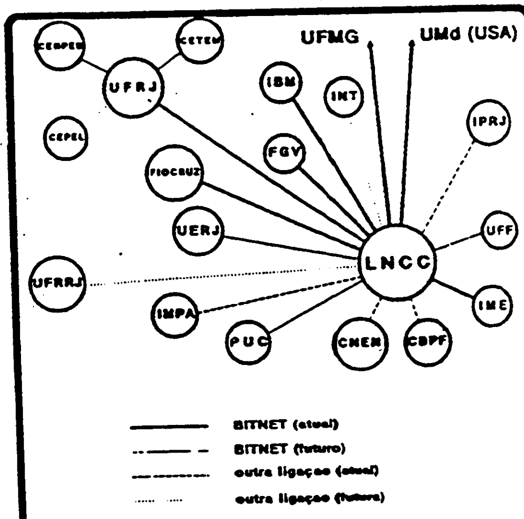

Fase 1 da Rede Regional RJ
(jan/91)

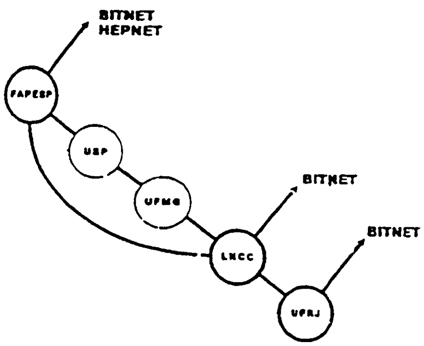

Conectividade Internacional
da RNP (jan/91)

III. MIGRATION TO INTERNET PROTOCOLS

02/91	FAPESP - FERMILAB (ESNET)
05/91 (?)	REGIONAL IP NETWORK IN RIO CONNECTED TO CERFNET AND SP.

IV. NATIONAL COORDINATION

CNPq BEGAN NATIONAL RESEARCH NETWORK (RNP) IN 07/90 WITH AIMS:

- o COORDINATION AND INTERCONNECTION OF REGIONAL NETWORKS
- o IMPROVEMENT OF INTERNAL COMMUNICATIONS INFRASTRUCTURE
- o COORDINATION OF INTERNATIONAL CONNECTIONS

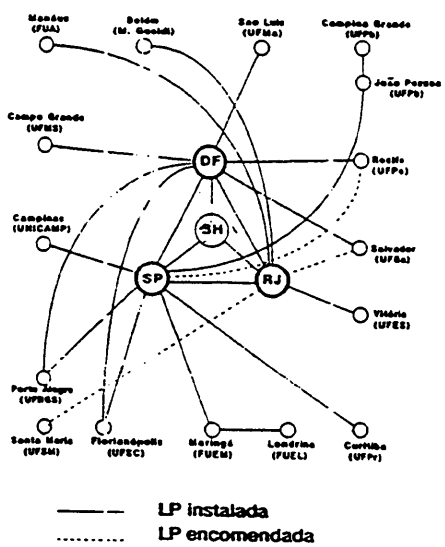

Integração da MECnet com a RNP

V. NEXT STEPS

1992
- 64 KBPS NATIONAL BACKBONE
- 64 KBPS CONNECTIONS
  BR - U.S.
  BR - EUROPE
- CLOSER LINKS WITH OTHER SOUT AMERICAN NETWORKS

4.2 DDN Mailbridge Report

Presented by Kathleen Huber/BBN

DDN MAILBRIDGES

Kathleen Huber

March 14, 1991

BBN Communications
A Division of Bolt, Beranek and Newman, Inc.

TOPICS

- Internet Growth
- DDN Mailbridges

INTERNET GROWTH SUMMARY

- 2136 Networks Advertised By BMILBBN on 3/5/91 - 2:13 PM
- 4648 Networks Registered

NUMBER OF NETWORKS

LINEAR
DECEMBER 1983-JANUARY, 1991

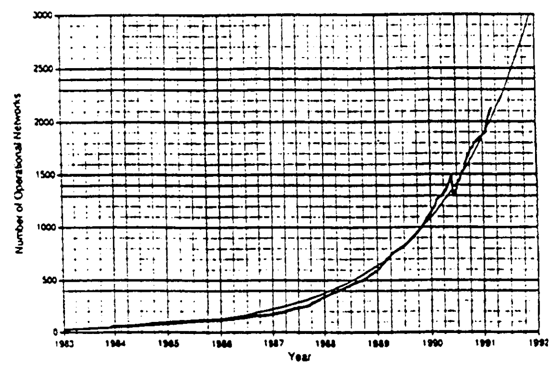

NUMBER OF NETWORKS

LOGARITHMIC
DECEMBER 1983-JANUARY, 1991

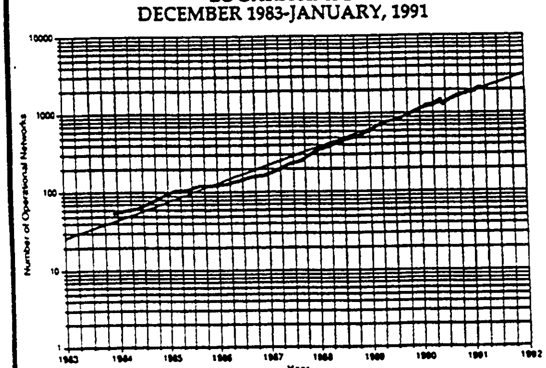

DDN MAILBRIDGES

EGP NEIGHBOR COMPARISON

	DIRECT NEIGHBORS								
	April 90'	June/July	Aug.	Sept.	Oct.	Nov.	Dec	Jan 91'	Feb
BMILAMES	90	68	55	52	54	50	58	60	61
BMILBBN	141	111	26	26	38	33	40	45	62
BMILDCEC	112	99	105	101	85	86	89	87	88
BMILISI	69	62	52	56	61	56	62	59	60
BMILOAK (BMILLBL)	43	63	72	73	70	72	59	62	78
BMILMTR	105	76	61	65	58	57	63	67	56
BMILRAN	-	-	33	30	42	26	38	40	50

TRAFFIC SUMMARY COMPARISON

	Avg. Pkts/Day Forwarded		Avg. Bytes/Pkts.		Avg. Pkts Dropped	
	Aug/Nov	Dec/Feb	Aug/Nov	Dec/Feb	Aug/Nov	Dec/Feb
BMILAMES	3,918,388	4,971,389	169	154	2.2%	0.7%
BMILBBN	148,521	173,642	463	460	0.6%	0.8%
BMILDCEC	324,852	350,733	307	245	8.6%	8.2%
BMILISI	90,910	291,263	406	342	0.9%	2.9%
BMILOAK (BMILLBL)	211,226	176,996	415	318	4.8%	6.1%
BMILMTR	1,905,353	2,091,064	178	182	0.4%	1.2%
BMILRAN	106,471	151,498	441	394	1.8%	2.4%

CURRENT STATUS

- Released Patch 9, "255 Gateways" Patch
  - Increase the Number of Networks the MAILBRIDGES are Capable of Supporting from 1800 Nets to 2700 Nets.
  - Increase the Number of Neighbors EGP is Capable of Supporting from 255 to 510; count of all external gateways in excess of 255 to be added to the count of internal gateways.
  - EGP Updates output by the "Split EGP" process will be throttled when buffers are not available for output.

CURRENT STATUS (cont.)

- Patch 10 Possibilities, "Token Queue" Patch
  - Correct the bug which contributed to token loss in buffer utilization.
  - Keep only one outstanding EGP peer Poll message on the Poll Queue thus decreasing the amount of polls to process and providing the EGP peer with only the most recent update.
  - Improve performance in customizing EGP Update messages.
- Load Balance Study
- Drop Rate Study

SUMMARY

- Current Efforts
  - Stabilize EGP and Store-and-Forward Service
  - Load Balance
  - Drop Rates
  - Prepare a Strategy for Long-term Growth Effects
- Future Possibilities
  - Begin Implementing a Strategy for Long-term Growth Effects
  - Synchronous Interface
  - Dynamically Manage Load Balancing
  - Decrease Drop Rates

4.3 ESnet Report

Presented by Tony Hain/LLNL

ESnet Status Report

IETF - Saint Louis

March 1991

Anthony L. Hain

Associate Network Manager

ESnet / NERSC

PAST ACTIVITIES:

Initiated Peering With CICnet

Added Connections To ORAU, Brazil

ER wide networking review

Deployed X.25 switching at FNAL, MIT, LLNL, BNL

STATS:

29 Routers Managed

71 Directly Connected Networks

698 Regional Connected Networks

614 Networks via Other Backbones

.7G Packets Reveived

70% IP / 30 % DECnet

PLANNED ACTIVITIES:

German 128kbps circuit...waiting DBP end...due 3/1/90

Interconnect With NSI DECnet @ FIX-E

Deploy cisco X.25 Switching at SLAC

Deploy Initial CLNP Routing Across Backbone

Start BGP Use On Development Net

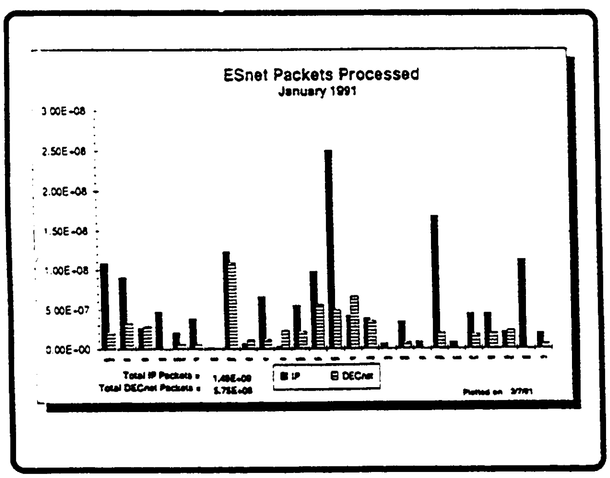

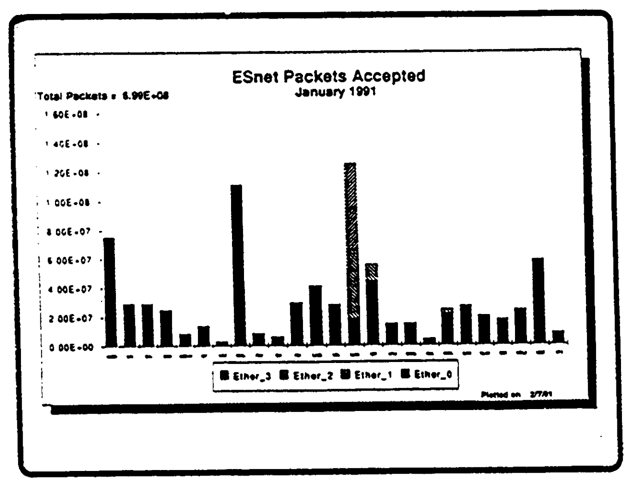

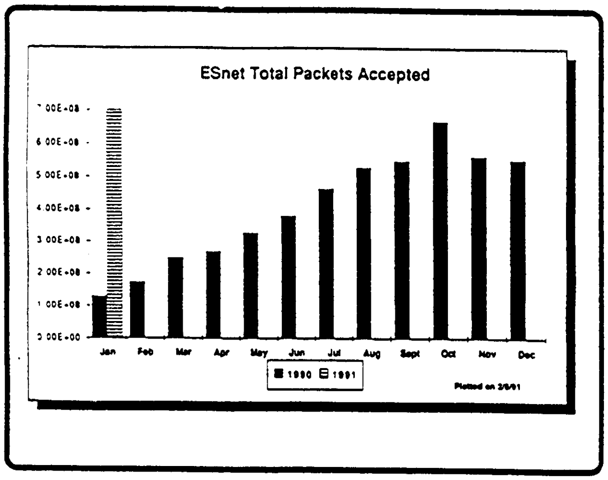
ESnet Total Packets Accepted
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sept
Oct
Nov
Dec
1990
1991

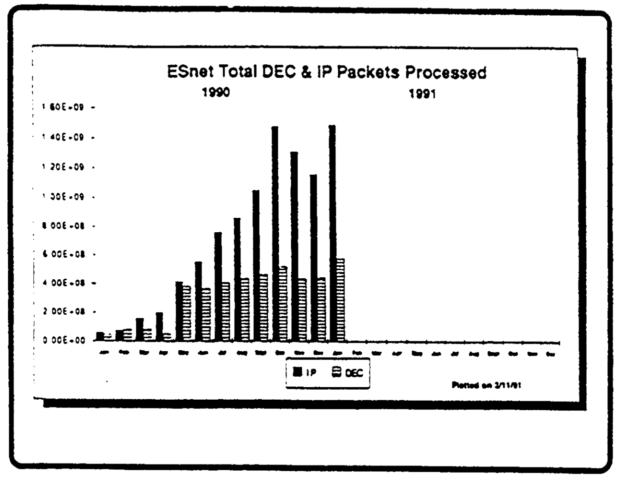
ESnet Total DEC & IP Packets Processed
1990
1991
IP
DEC

4.4 NSFnet Report

Presented by Elise Gerich/Merit

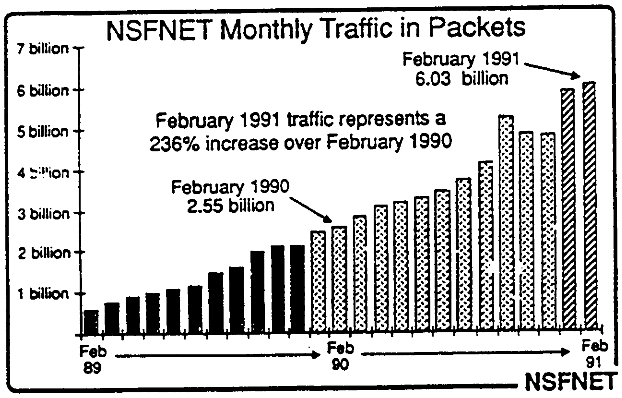
NSFNET Monthly Traffic in Packets
7 billion
6 billion
5 billion
3 billion
2 billion
1 billion
February 1991
6.03 billion
February 1991 traffic represents a
236% increase over February 1990
February 1990
2.55 billion
Feb 89
Feb 90
Feb 91
NSFNET

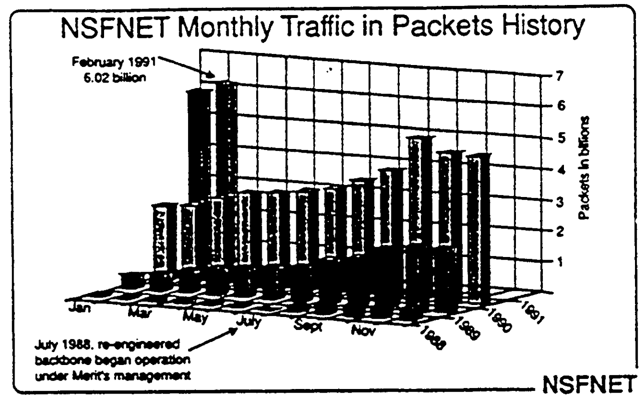
NSFNET Monthly Traffic in Packets History
February 1991
6.02 billion
Packets in billions
Jan
Mar
May
July
Sept
Nov
1988
1989
1990
1991
July 1988, re-engineered
backbone began operation
under Merit's management
NSFNET

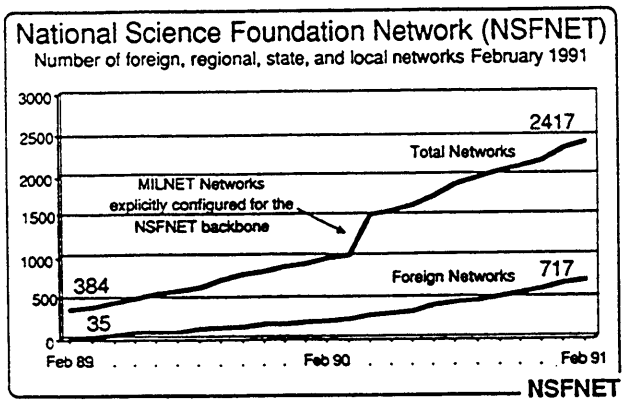
National Science Foundation Network (NSFNET)
Number of foreign, regional, state, and local networks February 1991
3000
2500
2000
1500
1000
500
0
2417
Total Networks
MILNET Networks
explicitly configured for the
NSFNET backbone
384
35
Foreign Networks
717
Feb 89
Feb 90
Feb 91
NSFNET

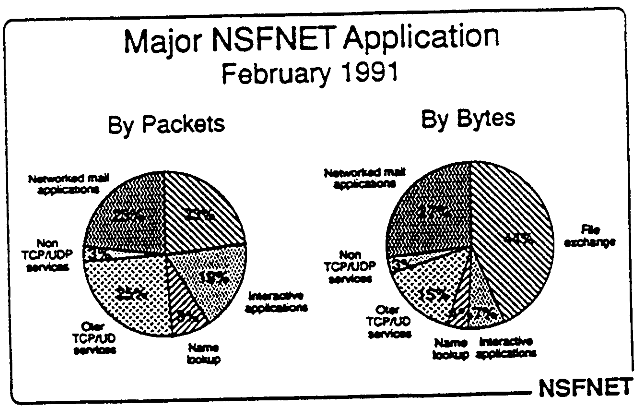
Major NSFNET Application
February 1991
By Packets
By Bytes
Networked mail applications
Non TCP/UDP services
Other TCP/UD services
Name lookup
Interactive applications
File exchange
44%
NSFNET

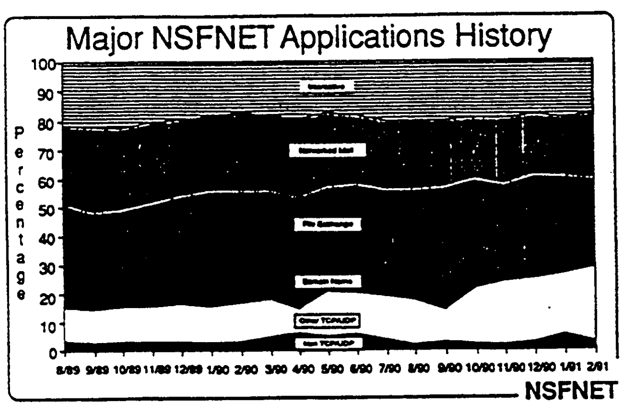
Major NSFNET Applications History
Percentage
100
90
80
70
60
50
40
30
20
10
0
NSFNET

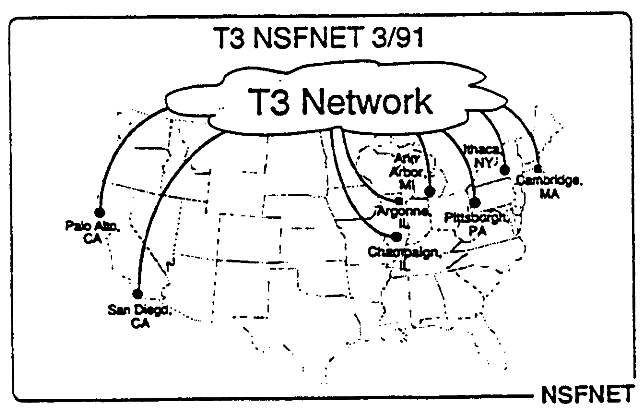
T3 NSFNET 3/91
T3 Network
Palo Alto, CA
San Diego, CA
Ann Arbor, MI
Argonne, IL
Champaign, IL
Ithaca, NY
Pittsburgh, PA
Cambridge, MA
NSFNET

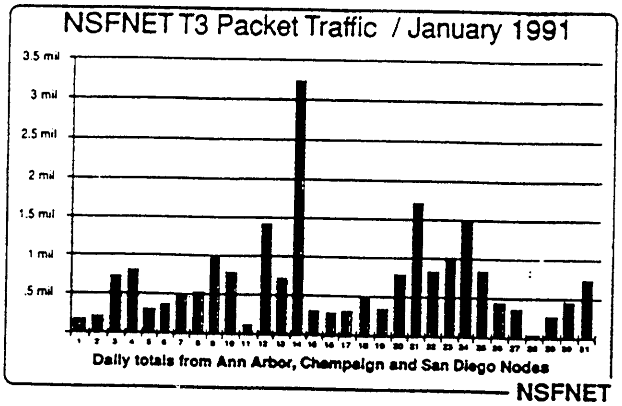
NSFNET T3 Packet Traffic / January 1991
3.5 mil
3 mil
2.5 mil
2 mil
1.5 mil
1 mil
.5 mil
Daily totals from Ann Arbor, Champaign and San Diego Nodes
NSFNET

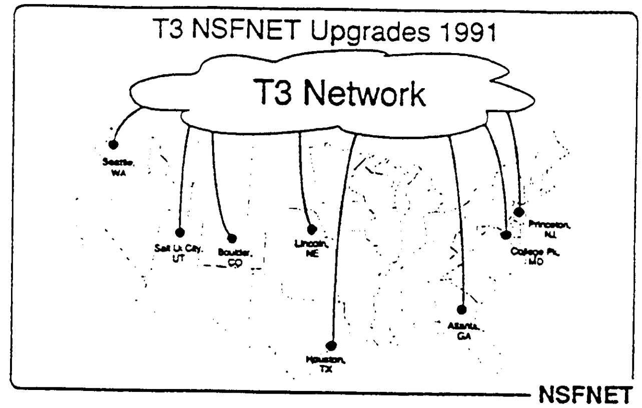
T3 NSFNET Upgrades 1991
T3 Network
Seattle, WA
Salt Lk City, UT
Boulder, CO
Lincoln, NE
Houston, TX
Atlanta, GA
Princeton, NJ
College Pk, MD
NSFNET

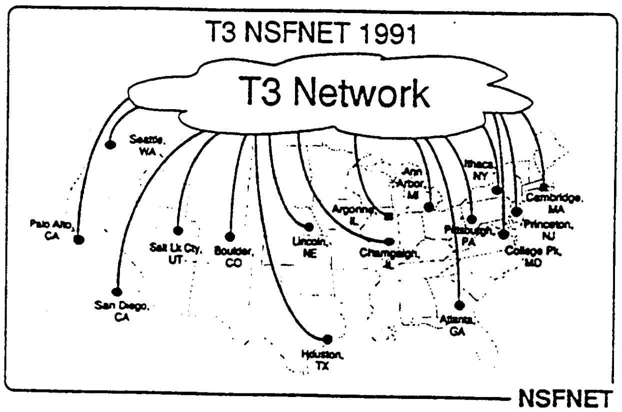
T3 NSFNET 1991
T3 Network
Seattle, WA
Palo Alto, CA
San Diego, CA
Salt Lk City, UT
Boulder, CO
Lincoln, NE
Houston, TX
Argonne, IL
Champaign, IL
Ann Arbor, MI
Ithaca, NY
Cambridge, MA
Pittsburgh, PA
Princeton, NJ
College Pk, MD
Atlanta, GA
NSFNET

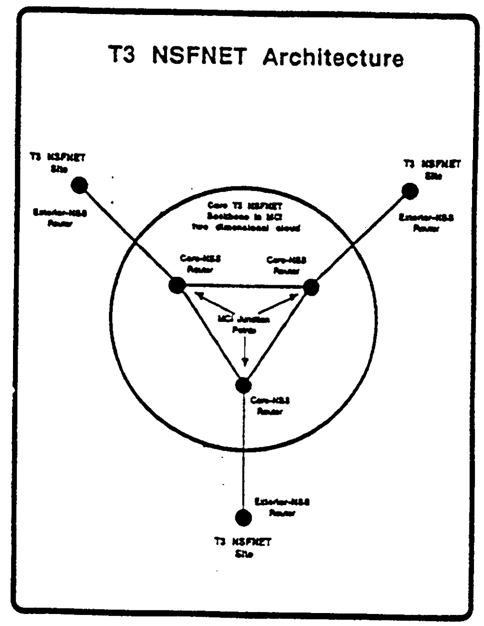
T3 NSFNET Architecture

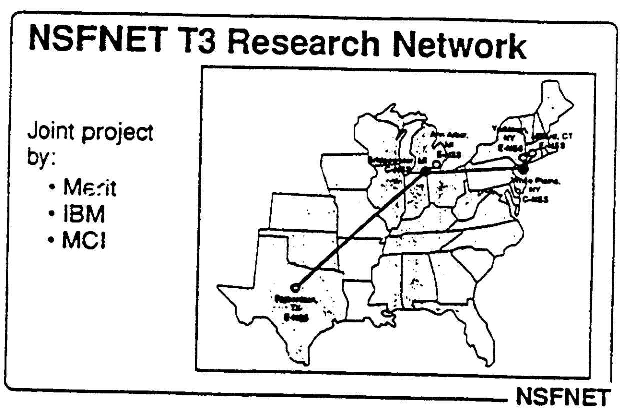
NSFNET T3 Research Network
Joint project by:
• Merit
• IBM
• MCI
NSFNET

4.5 NSI Report

Presented by Jeff Burgan/NASA

NASA Science Internet Status Report

March 7, 1991
IETF Meeting, St. Louis

Jeffrey G. Burgan
Sterling Software
NASA Science Internet Project Office
NASA Ames Research Center

Network Configuration

- 42 connected sites using Proteon routers
  - 23 running DECnet Phase IV in conjunction with TCP/IP
- 3 connections to the NSFnet
  - FIX-East (SURAnet, College Park, MD)
  - FIX-West (NASA Ames)
  - Rice University
- International Links
  - Australia (128K - being upgraded to 256K)
  - New Zealand (64K)
  - ISR, Japan (64K)

NASA Science Network

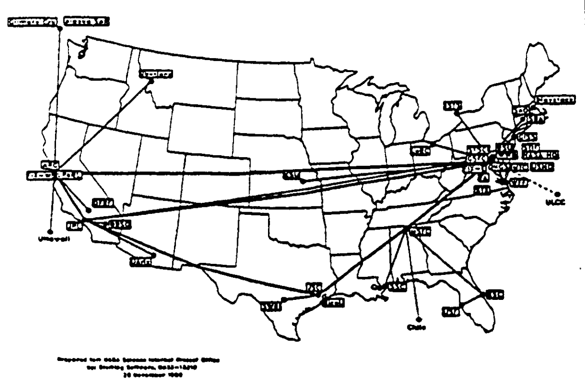

- Routing of DECnet Phase IV traffic over backbone links between ARC, GSFC and JPL
- INOC completed and operational
  - Network Operations Center supports both the NSI Wide Area Network and the ARCLAN Local Area Network

OSPF Update

- 15 routers running OSPF
  - 14 as part of the Backbone area
  - 1 in a Stub area
- Conversion to OSPF Version 2 on January 1, 1991
  - Stub area support
  - Inclusion of a forwarding address in an External LSA

NASA Science Internet OSPF System

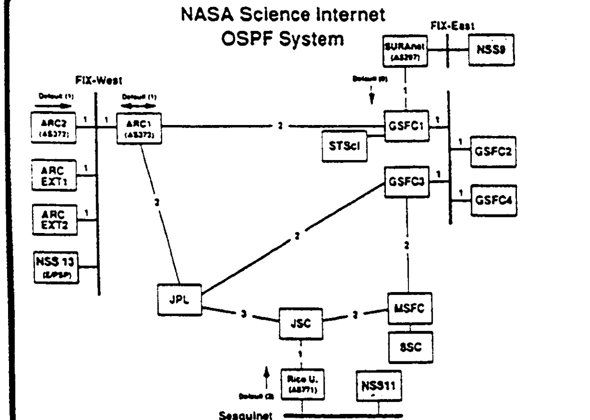

4.6 NORDUnet Report

Presented by Bernhard Stockman/NORDUnet

The European Situation ①

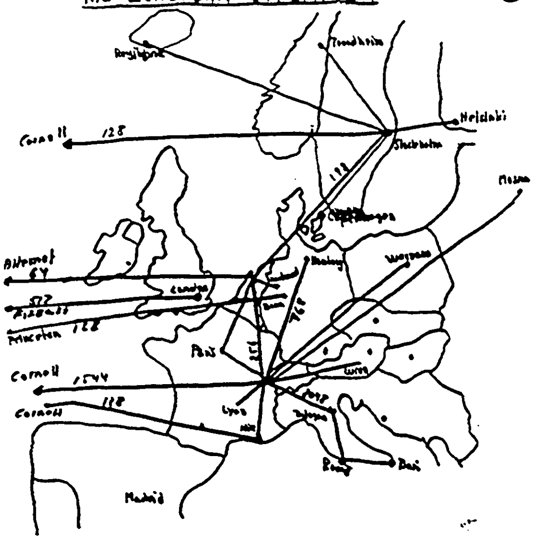

Some parts of the European infrastructure
(First) "Neutral Ethernet" in Amsterdam

European Engineering & Planning Group ②
(EEPG)

1. Estimation of traffic flows
2. Inventory of current infrastructure
3. Propose pan-European network technology
4. Suggest organizational model

First Report Planned
To Be Presented At
RARE Workshop in May 1991.

~ 5 M$ spent on international intra-European connectivity/year
+ 3 M$ on IXI/year

~ 5 M$ spent on transatlantic connectivity/year

Reseaux IP Europeens (RIPE) ③

- Coordination of IP-networking in Europe
- Accepted as a special group within RARE.
- Most European national and international IP-networks members
- Start up of RIPE Network Coordination Centre (NCC)
- RIPE whois database covering networks and persons at nic.eu.net (routers, AS's and lines proposed)
- ~ 40000 DNS registered hosts
- Topology maps on nic.eu.net as well as all kind of documents
- Increasing demand from East European countries on IP connectivity

Chapter 5

IETF Protocol Presentations

5.1 OSPF Status Report

Presented by Jeff Burgan/NASA, Rob Coltun/UMD and John Moy/Proteon

Operational data

Statistic	BARRNet	NSI	OARNet
Time interval	99 hours	277 hours	28 hours
Dijkstra frequency	50 min.	25 min.	13 min.
Ext. incrementals	1.2 min	.98 min	not gathered
Database turnover	29.7 min	30.9 min	28.2 min
LSAs per packet	3.38	3.16	2.99
Flooding rxmts	1.3%	1.4%	.7%

Link Bandwidth

- LSA flooded every 30 minutes
- 3 LSAs per packet
- per LSA: 23/100 bits per second
- Limiting to 5% of available bandwith:

Line Speed	# ext LSAs
9.6Kb	2087
56Kb	12,174

Router memory

- Approximately 64 bytes per LSA
- P4200 can store 10,000
- ways to reduce memory requirements:
  - Use default routing
  - Use stub areas
  - Use smarter EGP
  - Get away from flat address space

Router CPU

- Dominated by Dijkstra time
- Order (n * log (n)): n is no. of routers
- Steve Deering calculated:
  - on 10 mips machine
  - 15 milliseconds for 200 routers
- Other observations:
  - Dijkstra runs infrequently (15 min)
  - MILNET has 230 nodes
- Can break into areas to reduce Dijkstra cost

Designated Router

- receives and sends more packets than other routers
- switching DRs requires time
  - reason for DR
- Over 50 routers have been simulated on a LAN
- 13 routers on LAN in last round of testing
- Still, DR may limit # routers on LAN
  - Don't see this operationally

Third Round: First configuration

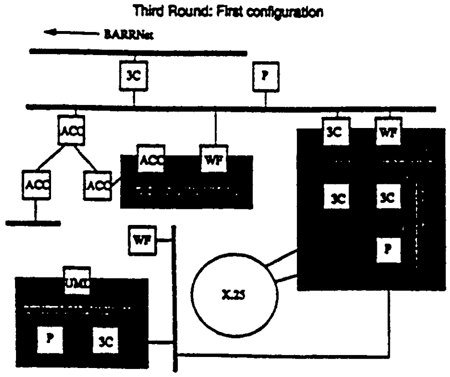

Figure 5: Third round testing: initial configuration

Third Round: Second configuration

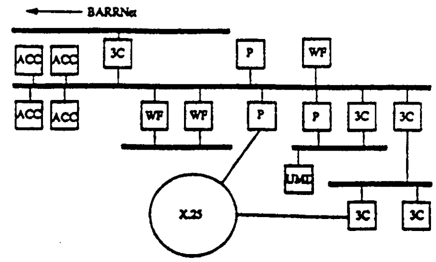

Figure 6: Third round testing: second configuration

Overall: Features tested

- Designated Router election
- Database Synchronization
- Flushing advertisements
- Import of external routing information
- Running over various network types
- Non-broadcast, multi-access networks
- Authentication
- Equal-cost multipath
- Variable-length subnet masks
- TOS routing

More features tested

- The following area features:
  - Extent of LSAs
  - Inter-area routing (summaries)
  - Virtual links
  - Stub areas

Simulation

- By Distributed Systems Research Group at University of Maryland baltimore County
- Using modified MIT simulator
- Using Rob Coltun's OSPF code
- Tested several different topologies, verifying
  - Database convergence
  - Area functionality
- Also formal analysis of:
  - Interface and neighbor FSM
  - DR election

Cost of protocol

- Link bandwidth
- Router memory
- CPU usage
- Role of Designated Router

Routing Overhead (RIP vs. OSPF)

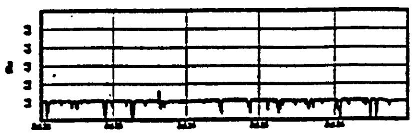

RIP (293 nets)

OSPF (394 nets)

Features tested operationally

- Designated router election
- Database synchronization
- Flushing advertisements
- Importing external routes
  - Type 1 and type 2
  - Forwarding addresses
  - External route tag
- Authentication
- Equal-cost multipath
- Stub areas

Deployment Limitations

- No multi-vendor deployments
- The following not used operationally:
  - Regular OSPF areas
  - Virtual links
  - Non-broadcast network (like X.25)
  - TOS routing

Interoperability testing

Site	Dates	#routers	#externals
Proteon	9/25/90	6	20-30
SURANet	12/17/90	10	96
3com	2/4/90	16	400

First round of testing

- 3com, Proteon and Wellfleet
- Tested:
  - Designated router Election
  - Database description process
  - Flooding
  - Importing external routes
  - Areas

Second round: Initial configuration

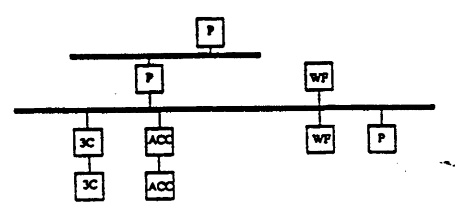

The NSI OSPF system

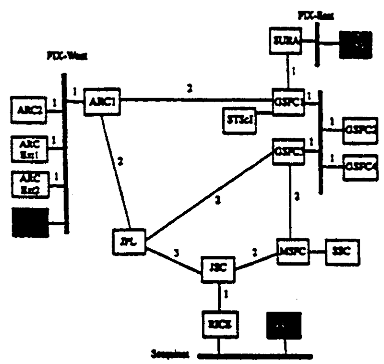

The BARRnet OSPF system

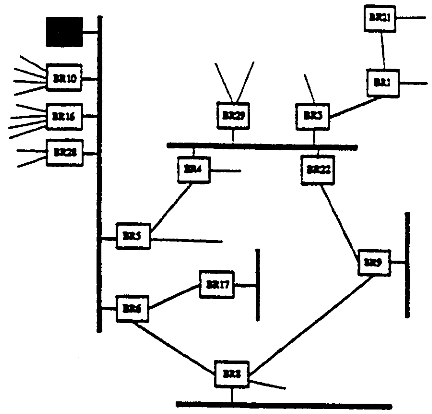

BARRNet Network Map

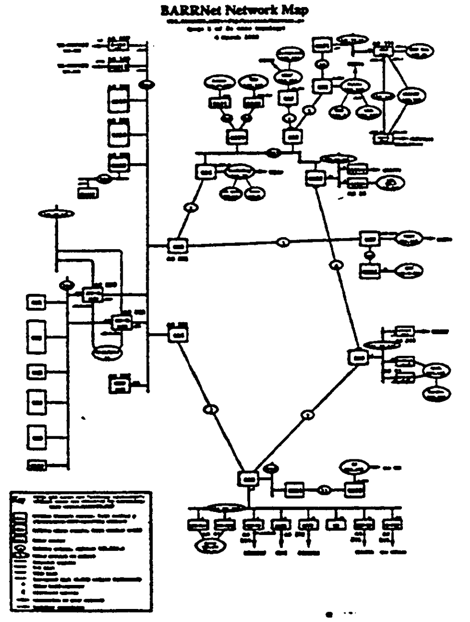

The OARnet OSPF system

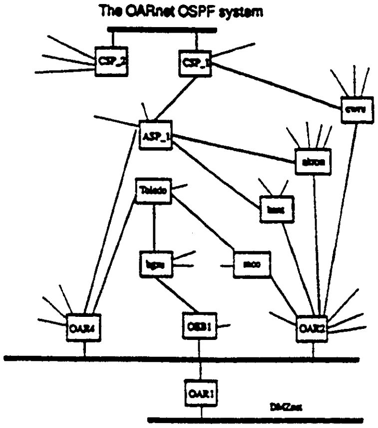

OSPF Route Convergence

```
PING east.gsfc.nasa.gov (128.183.104.4): 56 data bytes
64 bytes from 128.183.104.4: icmp_seq=0 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=1 time=140 ms
64 bytes from 128.183.104.4: icmp_seq=2 time=100 ms
64 bytes from 128.183.104.4: icmp_seq=3 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=4 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=5 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=6 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=7 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=8 time=80 ms

---- Link disabled ----

64 bytes from 128.183.104.4: icmp_seq=11 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=12 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=13 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=14 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=15 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=16 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=17 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=18 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=19 time=80 ms
64 bytes from 128.183.104.4: icmp_seq=20 time=80 ms
```

OSPF Status Report

- Documentation
- Management Information Base
- Security architecture
- Implementations
- Operational experience
- Interoperability Testing
- Simulation
- Performance Evaluation

Documentation

- V2 Specification published as an Internet Draft (1/91)
- Supersedes RFC 1131 (Version 1)
- Working Group mailing list: ospf@trantor.umd.edu
- OSPF testing mailing list: ospf-tests@seka.cso.uiuc.edu
- OSPF V2 MIB published as Internet Draft (3/91)

OSPF MIB

- Internet Draft (Rob Coltun and Fred Baker)
- SMI experimental code 23
- 127 objects overall
- 1 Group and 11 Tables
  - 18 Counters
  - 13 Gauges
  - 73 variables mandated by OSPF spec

Security Architecture

- All packet exchanges authenticated
- OSPF packet header contains
  - 16-bits of authentication type
  - 64-bits of authentication data
- Authentication types:
  - 0: Null
  - 1: Simple Password
  - 2-255: Reserved
  - > 255: Local use

Implementations

- Those participating in interoperability testing:
  - 3com
  - ACC
  - Proteon
  - Wellfleet
  - University of Maryland

Operational Experience

Site	V1 date	V2 date	RTs	Exts
NSI	4/90	1/91	15	496
BARRNet	4/90	11/90	14	1816
OARnet	10/90	not yet	13	140

5.2 BGP Status Report

Presented by Dennis Ferguson/U-Toronto, Sue Hares/Merit and Yakov Rekhter/IBM

BGP Status Report

Yakov Rekhter
T.J. Watson Research Center
IBM Corporation
e-mail:yakov@ibm.com

March, 1991

Design Objectives.

- Designed to address pressing problems in the Inter-AS routing.
    - Solve problems associated with EGP-2 within the existing topology.
    - Provide support for unrestricted topologies.
    - Inter-AS routing protocol.
- Flexible and expandable solution for the Inter-AS routing in TCP/IP Internets.
- Simple.

Protocol History.

- IWG/BGP working group (Guy Almes, Rice University, Yakov Rekhter, IBM).
- BGP-1 (RFC1105, June 1989).
- BGP-2 - Proposed Internet Standard (RFC1163, RFC1164, June 1990).
- BGP-3 (Internet-draft, January 1991).
- Implementation experience ⟺ refined specifications.

Protocol overview.

- Routing protocol with no restrictions on topology.
- Incremental updates ⇒ efficiency.
- Neither Link State, nor Distance vector.
- Efficient mechanism for routing loop suppression ⇒ fast convergence.
- Path attributes for route selection/routing policies.
    - Mandatory versus optional attributes ⇒ flexibility.
    - Introducing new attributes ⇒ expandability.
- Scaling O(networks).
- Security architecture - from simple authentication to encrypted digital signature.

Independent implementations.

- cisco Systems - Kirk Lougheed.
- gated - Jeff Honig, Dennis Ferguson.
- NSFNET Backbone - Yakov Rekhter.
- Specification to help implementors.
    - BGP-1 implemented in a 1 month (from scratch).
    - BGP-1 to BGP-2 - a week.
    - BGP-2 to BGP-3 - a day.

MIB for BGP.

- BGP-2 MIB (September 1990) - Steve Willis, John Burruss (Wellfleet).
- BGP-3 MIB (March 1991) - Steve Willis, John Burrus (Wellfleet).

Router Requirements and BGP.

- BGP Section - John Moy (Proteon), Jessica Yu (MERIT), Yakov Rekhter (IBM).

Operational Experience.

- NSFNET T1 Backbone since November 1989 (external BGP, phasing in internal BGP).
- CA*Net since the June 1990 (both internal and external BGP).
- NSFNET T3 Backbone since January 1991 (both internal and external BGP).
- Carries full complement of exterior routes (> 2000 networks).
- Graceful migration from EGP-2.

Conclusions.

- Good inter-AS routing protocols.
- Design based on proven technology.
- Addresses majority of the requirements.
- Practical to implement and use.

NSFNET Testing of BGP

Jessica Yu
Susan Hares
Merit/NSFNET Internet Engineering

NSFNET

NSFNET Testing of BGP

- Stages of BGP testing
- NSFNET T1-T3 interaction testing
- NSFNET T1 Experiences
- What next for NSFNET and BGP?

NSFNET

Stages of BGP testing

- NSFNET T1-T3 interaction testing
- NSFNET T1 Experiences
- What next for NSFNET and BGP?

NSFNET

Baby testing - Prototype testing

Can a protocol learn to walk and talk?

3 implementations tested out
rough first implementations

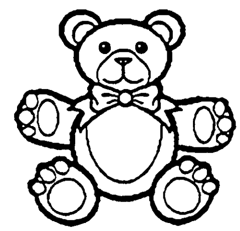

NSFNET

Initial Protocol Testing

Scenario -- cisco and gated machine connect to NSFNET Research Net

Tests -- can they run with out crash?
can they talk to each other (interoperable?)
do they send correct routing information ?
do they handle error situation correctly ?
....

Process -- debugging $\underset{\longrightarrow}{\overset{\text{n times}}{\longleftarrow}}$ bug-fixing

Results -- all worked satisfactory

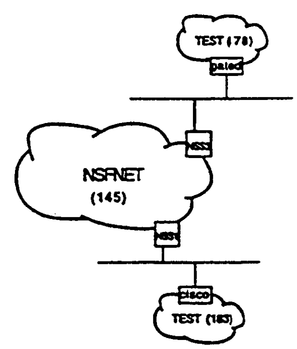

BGP Testing on the Research Network

Child testing - Conformance testing

Can a protocol learn how to speak and behave properly?

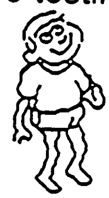

Can the implementations:

- send appropriate messages in the correct formats
- obeys the state machine described in the protocol
- are the normal error conditions handled (can the child be corrected?)

NSFNET

Teenager - First "Date" / Online Trials

- BGP implementations tried out at BARRNet, Merit, CICNet, EASInet, CA*Net
- BGP with EGP Back-up
- Internal BGP on T1 Test Network
- second phase on-line for almost a year

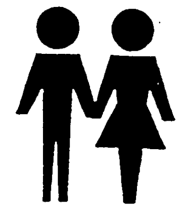

NSFNET

Further Protocol Testing

Scenarios -- Simulate the real network situation to verify BGP features

Tests -- does BGP do loop suppression ?
do the routes have the correct AS path ?
can policy be set to favor/disfavor certain pathes ?
does BGP really do incremental update as claimed ?
does the protcol handles error situation as described in spec?
does it do AS based filtering?

Further Protocol Testing (cont'd)

Results -- loop suppression -- yes
correct AS path -- yes
path selection -- routes originated from IGP favored
routes with shorter AS path favored
routes whose AS had more weight favored
incremental update -- indeed
AS based route filtering -- yes

AS x,y,z are BGP peer with each other
The two BGs within AS z also doing Internal BGP

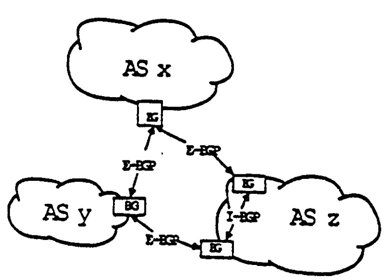

BGP testing simulating real Network

Application Testing

Scenario -- Simulating regional/NSFNET BGP connection using Research NSFNET

Tests -- IGP and BGP interaction at stub AS
Switch to backup routes
Coverging time
Router CPU utilization compare to EGP -- *reduced*
Traffic between the peer compare to EGP -- *goes down*
Routes recovery from outage compare to EGP -- *much fast*

NSFNET Backbone/regional BGP Implementation

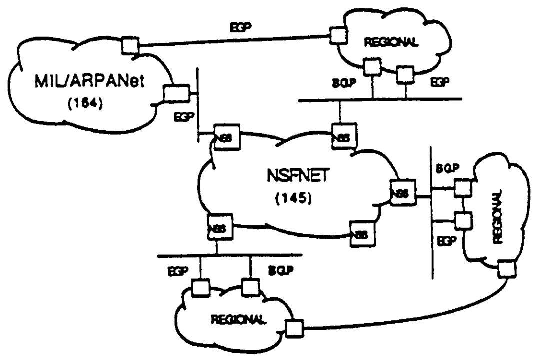

College Courses

- NSFNET T1 - NSFNET T3 Interconnect
- CA*Net
- T1 implementation

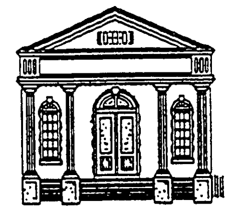

NSFNET

Deployment

Deploying BGP connection between NSFNET backbone and Regional Networks

Testing phases:

1t	cisco/gated BGP peer with NSS1 on the Research Network as the current EGP replacement
1.5t	cisco/gated do the filtering based on RD/AS number
2t	both peers do the filtering based on certain craterias
3t	all the NSSs on the RNet run BGP, without running Intra-RD/AS BGP
4t	all the NSSs on the RNet run BGP, with Intra-RD/AS BGP
5t	application of BGP policy functions

RNet -- the NSFNET Research Network

Deployment phases:

1d	volunteering sites BGP peer with the colocated NSS (EGP replacement)
2d	volunteering sites BGP peer with the colocated NSS, filtering based on thier craterias
3d	all the NSSs on the operational backbone run BGP code without running Intra-RD/AS BGP
4d	partial NSSs on the operational backbone run Intra-RD/AS BGP
5d	all the NSSs on the operational backbone run Intra-RD/AS BGP
6d	application of BGP policy functions

NSFNET T1 - T3 Interconnect

- Crash Courses
- Methods of "Getting" NSSs "down"
- Configuration Errors

NSFNET

Crash Courses

- Hard break
  - by circuit disconnect
  - by interface disconnect
- Software break
  - by software disconnect

NSFNET

Methods of "Getting" NSSs "down"

- Slow...one by one
- Taking down several NSSs at once
- Taking several NSSs up and down at once
- "MAD Monkey"

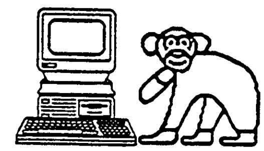

NSFNET

Configuration Error tested

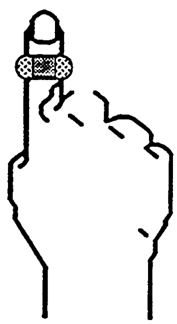

NSFNET

NSFNET T1 Experience

- Phasing over to Internal BGP in NSFNET T1
- EGP exists with BGP
- Routing Tables
- BGP is incremental
- Configuration Errors Again!!
- Traffic and CPU

NSFNET

Phasing Over Internal BGP in T1

- Two types of routing daemons
  - — IS-IS / IBGP
  - — IBGP only
- EGP exists with BGP
- Routing Table Behavior

NSFNET

Phasing Over Internal BGP in T1

- Incremental nature of BGP
- Configuration Errors Explored
- Reduction in traffic and processing time

NSFNET

What next for NSFNET and BGP?

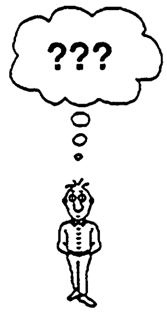

NSFNET

What next for NSFNET and BGP?

- Internal BGP into all of the T1 network
- NSFNET T3 uses Internal BGP
- Interconnect NSFNET T1 and T3 using External BGP

NSFNET

College Age protocol

Time to give the protocol a full-time job.
We'll learn and improve BGP as it starts to work full time, but EGP is approaching retirement.

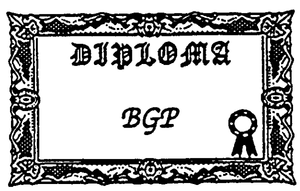

NSFNET

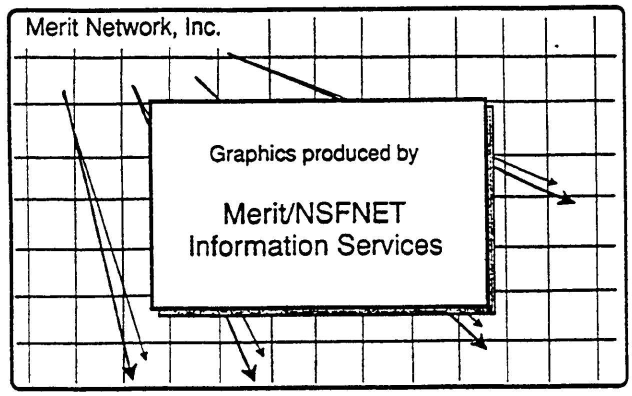

5.3 Alert-Management Status Report

Presented by Lou Steinberg/IBM

Status Report:

ALERT-MANAGEMENT Working Group

Chair:

Lou Steinberg
Louiss@ibm.com

Alert Management (Alert-Man) WG

1) CHARTER:

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 1098) and CMOT (RFC 1095).

2) OBJECTIVES:

A) 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.

B) Write an RFC detailing the above, including examples of its conformant use with both SNMP traps and CMOT events.

C) 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; examples of such techniques might include limiting number of alerts per time period, delayed reporting of "good news" (as in the link up sgmp trap on NSFNET), or the use of thresholds.

D) 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".

ALERT-MAN WORKING GROUP

1) STATUS
- *A) Documents*
  - *1) Flow Control Document Submitted to RFC* Bob
  - *2) lack of interest in defined alert RFC*
- *B) Delay in Document 1 publication*
- *C) Working Group Objectives Complete*

2) Flow Control Overview
- *A) Need*
- *B) Feedback*
- *C) Polled, Logged alerts*
- *D) MIB impact*
- *E) SNMP impact*
- *F) NSFNET implementation/test*

3) Final Thanks

DOCUMENT 1 (FLOW CONTROL) DELAY

- RFC - READY 3/90
- PUBLICATION AS PROPOSED STANDARD OR EXPERIMENTAL
  - 5+ IMPLEMENTATIONS EXISTED
  - HEAVY SUPPORT FROM VENDOR/USER COMMUNITIES
  - AREA DIRECTOR PLANNED TO SUBMIT AS PROPOSED
- MID 90 SHIFT: "EXPERIMENTAL"
- WG CONCERNS
  - "EXPERIMENTAL" IS NEW, + NOT ON STANDARDS TRACK
  - DEMAND + PROVEN UTILITY SAID NOT WORTH DOING A NON-STANDARD.

RESULT: • DOCUMENT SAT AS A DRAFT
• INTEREST WANED

DOCUMENT 1 DELAY:

TODAY:

DECISION TO PUBLISH AT THE LEVEL IAB FEELS IS APPROPRIATE

- EXPERIMENTAL
  - PUBLISH FOR HISTORICAL REASONS
  - FURTHER WORK UNLIKELY
- PROPOSED
  - SEE IF CRITICAL MASS EXISTS FOR NEW (RENEWED) DEMAND ... NEW IMPLEMENTATIONS

CLOSURE:

- DOCUMENT SUBMITTED TO PHILL GROSS TO:
  a) GET NEEDED MIB SUBTREE
  b) FORWARD TO RFC EDITOR AS EXPERIMENTAL OR PROPOSED STANDARD
- CLOSE WORKING GROUP ACTIVITIES

FLOW CONTROL - NEED

POLLING

- RELIABLE DETECTION
- REDUCE AGENT COMPLEXITY + PERFORMANCE HIT
- HIGH DETECT / RESPONSE TIME
- HIGHER TRAFFIC VOLUME (IMPACT INCREASES TO LOWER DETECT TIME)

ALERTS

+ LOW TIME TO DETECT PROBLEM
+ LOWER NETWORK TRAFFIC DEMAND WHEN OK.
- UNDER DELIVERY *
- OVER DELIVERY *

* SOLVE THESE

FEEDBACK

ASSUMPTIONS

- ALERTS CAN BE OVER PRODUCED (OPEN LOOP)
- THIS CAN ADVERSELY AFFECT NET + MANAGER
- ALERTS ARE PRIORITY, BUT NOT CRITICAL PATH (LOW COST TO INSTRUMENT)

SOLN

- COUNT ALERTS, LOG TIME.
- IF TOO MANY ALERTS PER TIME BASE; NOTIFY MANAGER AND DISABLE REPORTING. REQUIRE MANAGER APPROVAL TO RESUME.

RESULT

- BOUNDED OPEN LOOP BECOMES CLOSED LOOP, BASED ON SLIDING WINDOW THRESHOLD
- NO IMPACT TO ALERT GENERATOR PERFORMANCE UNLESS EXCEPTION CONDITIONS EXIST

DETAIL DISCUSSION

Polled, Logged Alerts

Assumptions

- Alerts can be under produced/delivered
- Polling is needed to detect this
- Need to minimize network traffic, agent burden, and detect time problems found in polling

Soln

- Local alert log gets polled
  - Fewer polls per agent, lower traffic and burden
  - One object polled per agent (or collection of agents) helps improve detect time.

Result

- Improves polling problems
- Scales well
- Retains polling strengths - reliable!

Detail Discussion

MIB / SNMP Impact

Feedback

maxAlertsPerTime	INTEGER
windowTime	INTEGER
alertsEnabled	INTEGER (Boolean)
alertsDisabled	ALERT (TRAP)

Polled, Logged

alertLog	SEQUENCE OF {logTableEntry}
logTableEntry	SEQUENCE { alertId INTEGER, alertData OPAQUE (ALERT) }
maxLogTableEntries	INTEGER

The above objects and trap are added to the MIB and protocol

Implementation Tested on NSFNET

Internal Ethernet monitor on: internal LANs and Telnet & 35.1.1 (Ann Arbor)

- SNMP based, with SNMP/SGMP-IP alerts
- Alerts on
  - Carrier lost
  - No carrier @ init
  - ARP conflict
  - H/W failure
  - Enet jam (unavailable)
    - Source known
    - Source unknown
  - ICMP source quench
  - Utilization threshold

→ It works!

- Proven
- Tested
- Useful

Final Thanks

- Alert Man Working Group is closing.
- Alert Man "Someday we really gotta get an RFC number". Informal subgroup plans one final meeting.

Chapter 6

Technical Presentations

6.1 University of Maryland Routing Simulator

Presented by Deepinder P. Sidhu/UMD

OSPF Simulation

University of Maryland
Baltimore County Campus
(UMBC)

20th IETF Meeting
St. Louis, MO
March 11-15, 1991

This research was supported, at the University of Maryland Baltimore County, in part by RADC and DARPA under contract number F30602-90-C-0010 to UMIACS at the University of Maryland. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official politics, either expressed or implied, of the Defence Advanced Research Agency, RADC, or the U.S. Government.

OSPF Project Contributors

→ Principal Investigator

Dr. Deepinder P. Sidhu
Professor of Computer Science
UMBC/UMIACS

→ Graduate Students

S. Abdallah
T. Fu
R. Nair

→ Rob Coltun

OSPF implementation

OSPF Features

→ OSPF is an Interior Gateway Protocol (IGP) for an Autonomous System (AS)

→ OSPF is a link state routing protocol

→ Splitting the AS into areas
- Contiguous networks and hosts grouped into areas
- Topology of an area invisible from outside
- Area isolation reduces routing traffic
- All routers within an area have identical databases
- Router can be in multiple areas
- Area border routers
- Stub areas

→ Backbone of the AS
- Special area
- Consists of networks not in any area
- Must be contiguous (may need virtual links)
- AS boundary routers

→ Supports several physical networks
- Point-to-point
- Broadcast
- Non-broadcast

OSPF Features

→ OSPF Routes
- Metrics
- Directed graph
- Cost on edges
- Topological database
- SPF run in each area
- Shortest-path tree
- External routing information (leaves)
- Equal-cost paths
- TOS-based routes
- IP subnetting
- Authentication of routing updates

→ Inter-area routing
- Area border routers inject information about routing outside the area
- Area border router is connected to backbone

→ AS external routes
- AS boundary routers flood information about external destination into AS

→ Partitions of areas
- No repair of area partition
- Partitioned area becomes 2 or more separate areas
- Backbone must not partition

OSPF Features

- OSPF routing packets
  - Hello
    - discover/maintain neighbors
  - Database description
    - summarizes database content
  - Link-state request
    - database download
  - Link-state update
    - database update
  - Link-state ACK
    - flooding acknowledgement
- Router adjacencies
  - Hello protocol establishes/maintains neighbor relationships
  - Router forms adjacencies with new neighbors
  - Routing packets are sent on and received only on adjacencies
  - Topological databases are synchronized between pairs of adjacencies
  - DR determines adjacencies on multi-access networks

OSPF Features

- Designated Router (DR) and Backup Designated Router (BDR)
  - For each multi-access network
  - Elected by Hello Protocol
  - DR end point of many adjacencies
  - DR sends link state update packets
  - BDR is backup to DR
- Flooding within an area
  - Reliable flooding of link state advertisements
  - Routes computed from exactly same topological database

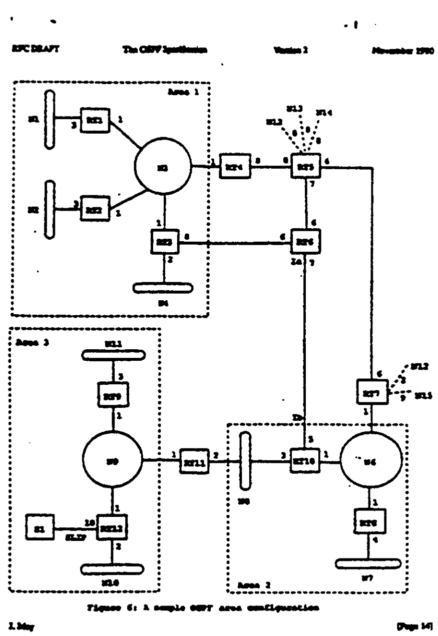

OSPF Features

- Neighbour State Machine

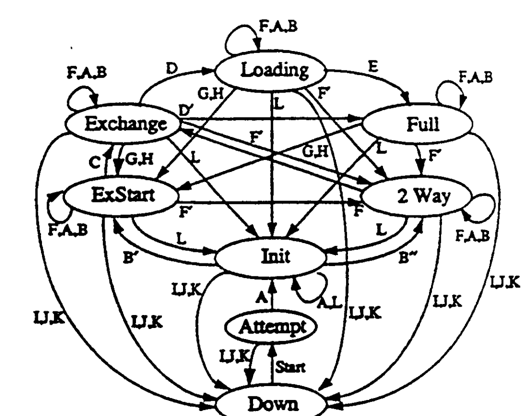

Events

- (A) - Hello Received
- (B) - 2 Way Received
- (B') - 2 Way Received [either present or the neighbouring] [router is DR/BDR]
- (B'') - 2 Way Received [neither is DR/BDR]
- (C) - Negotiation done to decide Master/Slave relation
- (D) - Exchange Done
- (K) - Inactivity
- (L) - 2 Way Received
- (E) - Loading Done
- (F) - Adj OK ? [Must form] [adjacency]
- (F') - Adj OK ? [Need not form] [adjacency]
- (G) - Sequence number mismatch
- (H) - Bad LS Req
- (I) - Kill Neighbour
- (J) - LL Down

OSPF Features

→ Interface State Machine - Multi-access Network

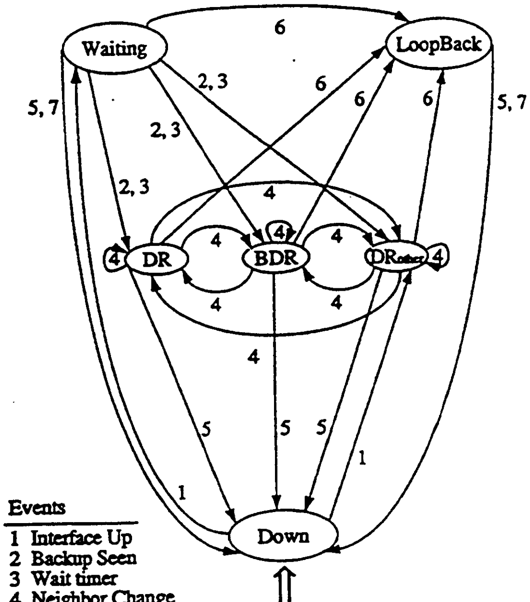

Events

1 Interface Up
2 Backup Seen
3 Wait timer
4 Neighbor Change
5 Interface Down
6 Loop Ind
7 Unloop Ind

Network Simulator

→ Event-driven
- Significantly enhanced version of MIT simulator

→ Modular design
- Easily modifiable

→ Interactive graphical interface
- X Window system

→ Simulator components
- Links
  - Point-to-point
  - Ethernet
- Switches (Routers)
- Hosts
- Data sources
  - Poisson distribution
  - Uniform distribution
  - Normal distribution
  - Exponential distribution
- Data sink

Network Simulator

→ User interface features

- Monitoring
  - Displays total number of packets in network
  - Displays average packet transit time
  - Displays average packet size
  - Displays convergence time
  - Logs routing table changes
  - Logs link state database changes
  - Logs packet information (type, size)
- Windowing
  - Displays a node and its neighbors
  - Displays complete topology
- Displaying data structures
  - Routing tables
  - Neighbor state machine
  - Interface state machine
- Tracking simulation in real time
  - Observation of meters (switch utilization, number of packets, queue lengths, memory utilization, ...)

Network Simulator

- Creating networks
  - Generates random topologies
- Creating topological changes in network
  - Triggers events using predefined scripts
  - Triggers events interactively
- Printing
  - Prints simulator screens
- Batch processing

OSPF Simulation

→ Preliminary Results

- Tested OSPF on
  - large networks (> 50 nodes)
  - networks with multiple areas
  - heterogeneous networks
    (point-to-point & broadcast)
- Checked correctness of link state databases and routing tables
- Experiments
  - convergence time
  - number of update messages

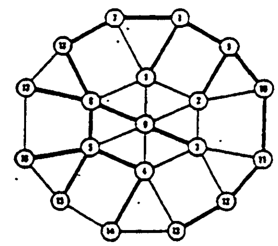

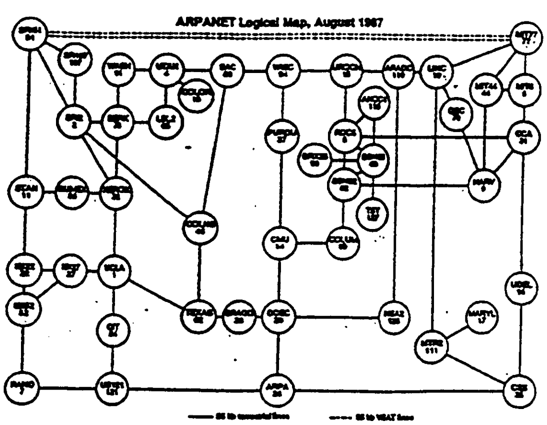

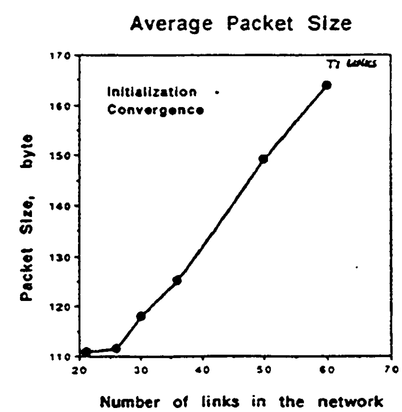

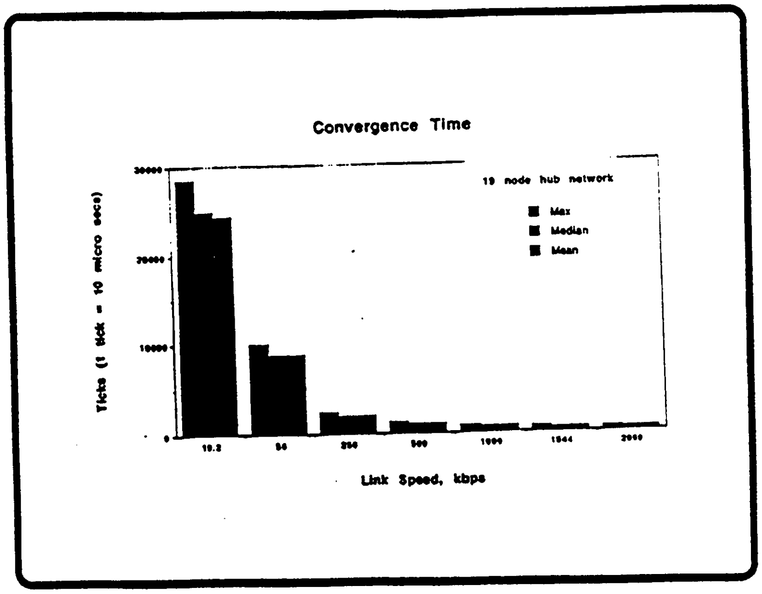
Convergence Time
19 node hub network
Max
Median
Mean
Ticks (1 tick = 10 micro secs)
Link Speed, kbps

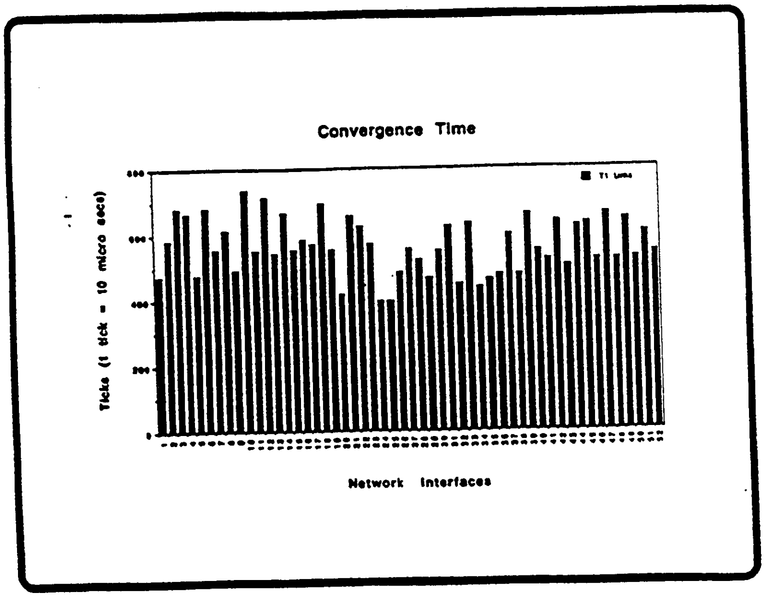
Convergence Time
Ticks (1 tick = 10 micro secs)
Network Interfaces

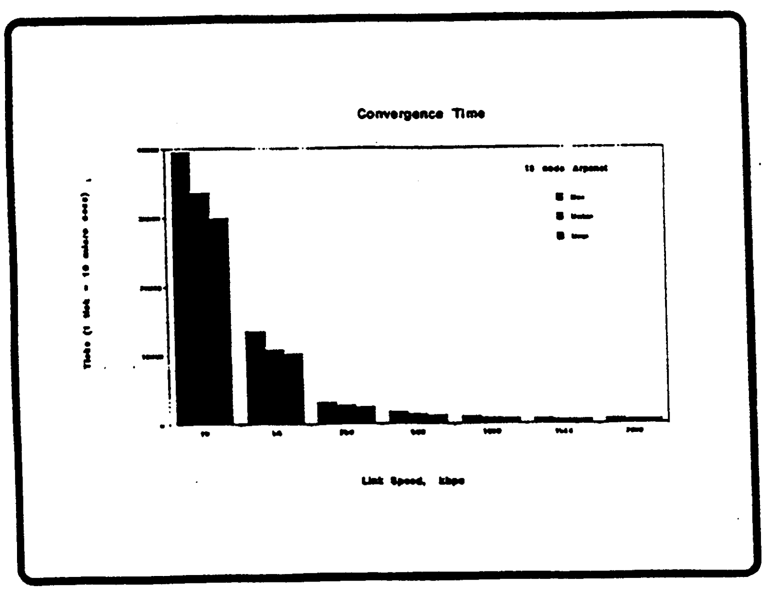
Convergence Time
Link Speed, kbps

OSPF Simulator

→ Analysis in Progress

- Performance measurements
  - quality of paths
  - number of update messages
  - convergence time
- Parameter changes
  - network size
  - link speed
  - offered load
- Test scenarios
  - single link failure/recovery
  - single node failure/recovery
  - multiple link & node failures
    - simultaneous failures
    - intermittent failures
    - random failures
    - chained sequence of failures
    - overlapped failure/recovery sequences
- Link state database synchronization
  - consistency of topological databases
- OSPF election
  - interface & neighbor state machines (deadlocks, incompleteness errors)
  - convergence
  - consistent view of DR & BDR

OSPF Simulator

→ Demos

- Demo 1: Simulator Features
- Demo 2: OSPF Databases
  - 12 node network (3 Areas)
  - Interface & neighbor state machines
    - neighbor discovery (Hello Protocol)
    - bidirectional communication
  - Internal router
    - Link state database
  - Area border router
    - Link state database
- Demo 3: OSPF Data Forwarding
  - Across multiple areas
- Demo 4: OSPF Electiion
  - 8 node network (broadcast & point-to-point)
  - 50 node broadcast network
  - Election in progress
  - Topological changes (DR or BDR or both go down)
- Demo 5: OSPF Backbone Partitioning
  - No repair at level 2
  - dropping of data packets
- Demo 6: OSPF & Large Networks
  - 50 node network (broadcast & pt-to-pt)

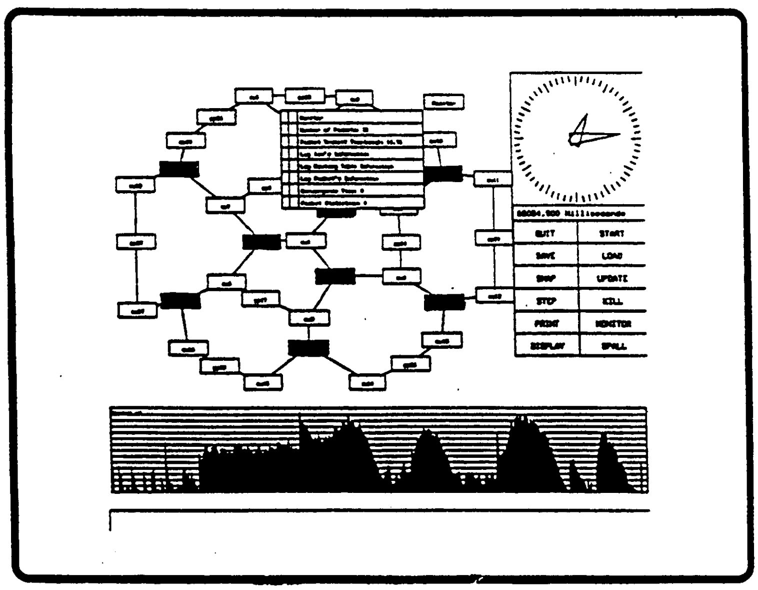

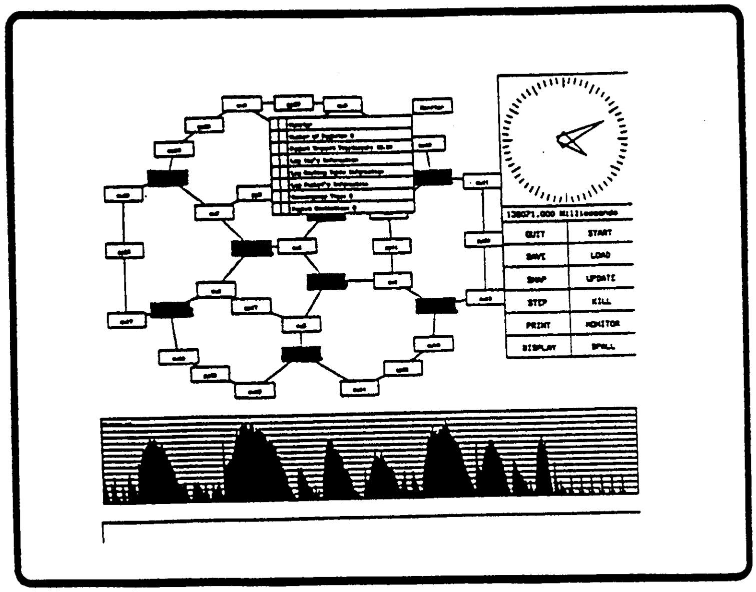
QUIT
START
SAVE
LOAD
SNAP
UPDATE
STEP
KILL
PRINT
MONITOR
DISPLAY
SPALL

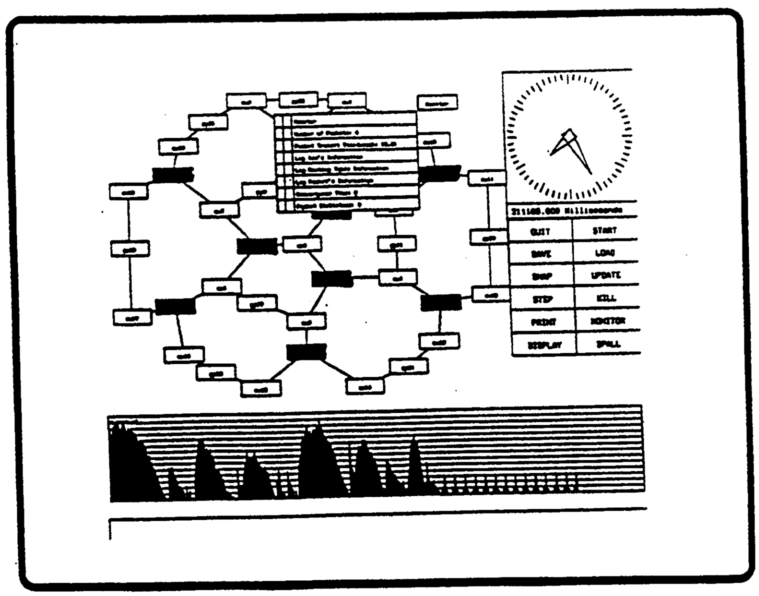
QUIT
START
SAVE
LOAD
SNAP
UPDATE
STEP
KILL
PRINT
MONITOR
DISPLAY
SPALL

6.2 Architecture and Goals for Interim Interagency NREN

Presented by William Johnston/LBL and Peter Ford/LANL

INTERAGENCY INTERIM

NATIONAL RESEARCH AND EDUCATION

NETWORK (NREN)

ARCHITECTURE AND IMPLEMENTATION PLAN

PRESENTED BY:

BILL JOHNSTON - LBL

PETER FORD - LANL

MARCH 12, 1991

IETF - ST. LOUIS

EMAIL: NREN@CNT.LANL.GOV

WHAT IS THE NREN?

- COMPONENT OF HPC
- EXPAND & ENHANCE U.S. PORTION OF INTERNET
- SUPPORT EDUCATION
- SUPPORT SCIENCE & RESEARCH COMMUNITIES
- SUPPORT BASIC RESEARCH & HUMAN RESOURCES PART OF HPCC

NREN WILL EVOLVE:

- FROM CURRENT INTERNET
- BY INCREASED COOPERATION WITH INDUSTRY
- BY APPLICATIONS OF LESSONS LEARNED FROM INTERNETWORKING RESEARCH
- NEW USER BASE
- NEW APPLICATIONS

NREN IS NOT AND END POINT !

HOW CAN WE MEASURE SUCCESS OF THE NREN?

- SUPPLIES BANDWIDTH AS REQUIRED BY SCIENCE & EDUCATION
- SUPPLIES CONNECTIVITY AS REQUIRED
- IMPROVES QUALITY OF EDUCATION
  - MORE QUALIFIED PROFESSIONALS
  - IMPROVEMENT IN CURRICULUM
- PROGRESS ON SOL'N OF GRAND CHALLENGES
- PEOPLE ARE WILLING TO PAY FOR IT.
  - DOES NOT IMPLY TOTAL PRIVATIZATION
- NETWORK CONNECTION $\not\Rightarrow$ NETWORK STAFF
- IETF MEETS ONCE A YEAR

NREN COMPONENTS IN THE FY92 PROGRAM:

- GIGABIT R&D LEAD AGENCY IS DARPA
- INTER-AGENCY INTERIM NREN

  LEAD AGENCY IS NSF

  DIRECTOR, NSF DIVISION OF NETWORKING AND COMM. RESEARCH AND INFRASTRUCTURE - STEPHEN WOLFF

  HIRES LANL/LBL TEAM TO DEVELOP AN ARCHITECTURE FOR INTER-AGENCY INTERIM NREN (JAN 91, DUE MARCH '91)

  PROJECT MANAGERS -

  JOHN MORRISON - LANL
  SANDY MEROLA - LBL

  TEAM -

BOB FINK	LBL	MITCH SUKALSKI	LANL
PETER FORD	LANL	ED THEIL	LBL
BILL JOHNSTON	LBL	C. PHIL WOOD	LANL

  TECHNICAL CONSULTANTS -

LARRY LANDWEBER	U. WISC.
VAN JACOBSON	LBL

E-MAIL: NREN@CNT.LANL.GOV

PROJECT PLAN

1) REQUIREMENTS
2) SUPPORTING TECHNOLOGIES
3) DIVE INTO ARCHITECTURE
4) IMPLEMENTATION PLAN
5) MANAGEMENT PLAN
6) DISCUSS @ IETF
7) REDO 1-5
8) REPORT TO WOLFF @ NSF

SUPPORTING TECHNOLOGY

- ON THE WAY TO GIGABITS / SEC
- TELECOMMUNICATION INDUSTRY STDS. (HERE OR COMING SOON)
  - FRAME RELAY
  - SMDS
  - B-ISDN / ATM
  - SONET
- T - 1 / T - 3 WILL STILL BE AROUND!
- ROUTERS
- IP / CLNP / ???
- ROUTING PROTOCOLS
- NIC / NOC SOFTWARE / PROTOCOLS

EVOLUTION

LATE 60s, EARLY 70s

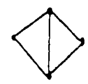

LANS TAKE OFF

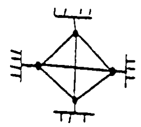

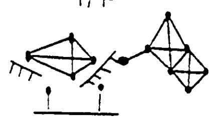

THINGS START GETTING INTERESTING!

WHAT IS RIGHT ABOUT CURRENT INTERNET

- HAS SUCCESSFULLY EVOLVED!
- IP INSURES INTEROPERABILITY WITH MEDIA IND.
- LOOKS "FREE" TO MANY END USERS
- NEW APPLICATIONS / PROTOCOLS
- ATTRACTS COMMUNITY THAT IS WILLING TO MAKE THE INTERNET WORK

PROBLEMS:

- SCALING
  - IP ADDRESS SPACE
  - ROUTING (SEMANTICS OF IP?)
- ACCOUNTING, LEADS TO LIMITED POLICY CHOICES
  - LIMITS INDUSTRIAL PARTICIPATION
- SECURITY - LITTLE AGREEMENT ON WHAT IT IS AND HOW TO ACCOMPLISH IT.
- REQUIRES NETWORK EXPERTISE TO CONNECT AS AN END NODE.
- RESOURCE GUARANTEES

TODAY: (ONE WAY OF LOOKING AT THINGS).

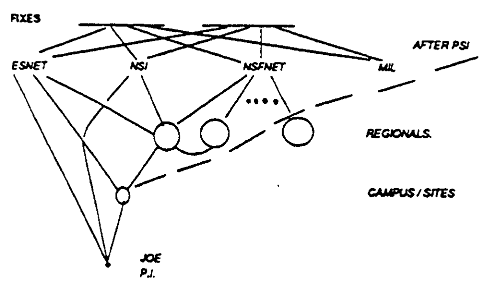

NOTE: THIS IS AN OVERSIMPLIFICATION
- NO INTERNATIONAL LINKS SHOWN

TWO MODELS

1) LEVEL 1 BIT PIPES (T - 1 / T - 3) WITH LEVEL 3 ROUTERS (MULTIPROTOCOLS)

- YIELDS A NETWORK OF NETWORKS
- COST BENEFITS:
  COST OF 8 - 10 T - 1's = COST OF 1 T - 3
- AVAILABLE TODAY, ESTABLISHED SERVICE
- PROBLEMS:
  STATIC DEMULTIPLEXING OF BANDWIDTH

2) LEVEL 2 PACKET ORIENTED SERVICE WITH LEVEL 3 ROUTING

- YIELDS COMMON SERVICE PROVIDERS

? COST BENEFITS - TOO EARLY TO TELL

? INDICATION THAT LEVEL 2 SERVICES WILL EMERGE IN FY92/FY93
  ADVANTAGE -LINK LEVEL RESOURCE GUARANTEES
  -INFRASTRUCTURE ⇒ MORE RELIABLE, MORE ROBUST NET.

? GEOGRAPHIC DISTRIBUTION

IMPLEMENTATION PLAN

- CONTINUE T - 1 ⟹ T - 3 IMPROVEMENTS (FAST MODEL)
- MOVE FROM MODEL 1 TO MODEL 2
  - COST CONSIDERATIONS
  - AS REQUIRED TO SOLVE HPCC PROBLEMS
  - STIMULATE WIDE GEOGRAPHIC COVERAGE WITH LEVEL 2 SERVICE
- NEED TO IMPLEMENT ACCOUNTING
- "GRAND CHALLENGES OF NETWORKING"
  - SCALING
  - ROUTING
  - SECURITY
  - RELIABILITY / ROBUSTNESS
  - END TO END USER SERVICES
    NIC / NOC / NSC

EVOLUTION TOWARDS FULL MESH!

REQUIREMENTS:

Requirements cover functionality, operating environment, service guarantees, and Federal program needs.

The NREN is a component of the Federal High Performance Computing and Communications program, and, as such, inherits a collection of programmatic goals and requirements. See, for example:

- Office of Science and Technology Policy, The Federal High Performance Computing Program, September 8, 1989.
- Federal Research Internet Coordinating Committee, Program Plan for the National Research and Education Network, May 23, 1989.
- Federal Networking Council (FNC) Engineering and Planning Group, Architectural Requirements for the National Research and Education Network, Version 1.0, January 22, 1991.
- Federal Coordinating Council for Science, Engineering, and Technology, Grand Challenges: High Performance Computing and Communications, Supplement to the President's Fiscal Year 1992 Budget.

William Johnston/Lawrence Berkeley Laboratory

In general these can be summarized as:

- Acceleration of the commercial availability and utilization of the next generation of high performance computers, and wide spread access to large data repositories, to extend U. S. leadership in these areas.
- Provide wide dissemination and application of these technologies to serve the national economy, national security, education, and the protection of the global environment; of high performance computers to extend U. S. leadership in these areas.
- Spur gains in U. S. productivity and industrial competitiveness by making these technologies part of the manufacturing design and production process.
- Support solution of important scientific and technical challenges.
- Support the underlying research, network, and computational infrastructure on which U. S. high performance computing technology is based.
- Support the U. S. human resource base in order to meet the needs of industry, universities, and government.

William Johnston/Lawrence Berkeley Laboratory

In terms of the NREN itself, we can enumerate the general requirements as:

- Provide a design and infrastructure for a high degree of end user connectivity over the entire geographic extent of the U. S.
- Work toward the cooperation with the telecommunications industry that will lead to the successful development of a broader, privately-operated national information infrastructure.
- Support the solution of scientific grand challenge problems by providing for the required connectivity between collaborating scientists, and the computing systems that they use.
- Provide service and connectivity guarantees to Federal Agencies independent of other, non-Federal networks, and of other users of the NREN.
- Promote the HPCC goals of uniform and ubiquitous access for the U. S. scientific and education communities.
- Establish a model for needed access security.
- Support in the short term, the establishment of "virtual private networks" (VPN) as a mechanism for service guarantees.
- Research and development on policy based routing mechanisms and security.
- Provide for the identification of needed accounting.
- Handle multiple protocols.
- Interconnect with international networks.
- Allows for substantial scalability.

William Johnston/Lawrence Berkeley Laboratory

- Provide for a collection of network services like domain name service, "white" and "yellow" pages like services, mail gateways to other networks and medium.
- Accommodate the addition of commercial information services.

William Johnston/Lawrence Berkeley Laboratory

MANAGEMENT:

Management structure in support of the NREN needs to ensure:

- Effective project management for IINREN implementation.
- User oversite.
- Technical oversight.
- Needed management structure is not yet determined.

William Johnston/Lawrence Berkeley Laboratory

NREN Management structure must:

- Enable the transition from NSFnet to Interim Interagency NREN to Gigabit NREN.
- Facilitate Federal agency liaison in NREN planning and coordination.
- Facilitate coordination of the Government, Educational, and Industrial participants.
- Provide for user community control over the network directions, and offerings.
- Provide for effective technical review.
- Identify and coordinate needed research.
- Promote telecommunications industry participation.
- Facilitate a phased transition to commercial support.
- Ensure effective interfacing with other networks.

William Johnston/Lawrence Berkeley Laboratory

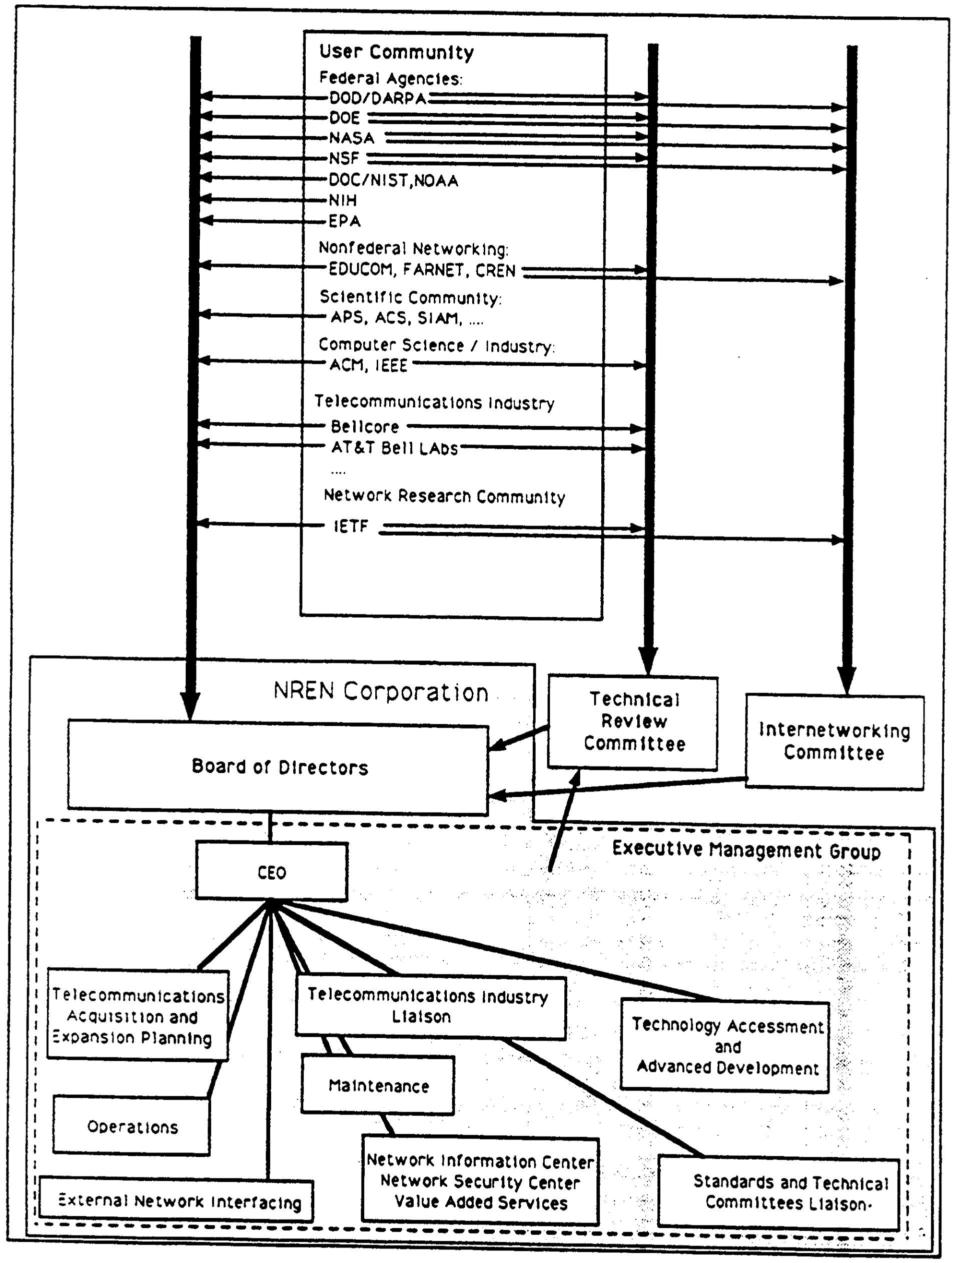
User Community
Federal Agencies:
DOD/DARPA
DOE
NASA
NSF
DOC/NIST,NOAA
NIH
EPA
Nonfederal Networking:
EDUCOM, FARNET, CREN
Scientific Community:
APS, ACS, SIAM, ...
Computer Science / Industry:
ACM, IEEE
Telecommunications Industry
Bellcore
AT&T Bell LAbs
....
Network Research Community
IETF
NREN Corporation
Technical Review Committee
Internetworking Committee
Board of Directors
Executive Management Group
CEO
Telecommunications Acquisition and Expansion Planning
Telecommunications Industry Liaison
Technology Accessment and Advanced Development
Maintenance
Operations
Network Information Center
Network Security Center
Value Added Services
External Network Interfacing
Standards and Technical Committees Liaison

6.3 Experiments DEC-bit Congestion Avoidance

Presented by Rick Wilder/MITRE

To support experimentation with the DEC-bit Congestion Avoidance scheme, the algorithms were implemented in the OSI Connectionless Network Protocol (CLNP) and TP4. The protocols developed at the University of Wisconsin and BSD 4.3 were used. Experiments were then run in an internet testbed network at the MITRE Corp. to measure the effects of the CA implementation with a variety of configurations and traffic loads to test the effectiveness and generality of the CA.

One set of experiments that was performed to test the applicability of CA to realistic network environments used intermittent as well as constant-demand traffic generators. The constant-demand traffic generators send data at the maximum rate that the network will accept it until a pre-determined amount is sent. The intermittent traffic generators start transmission bursts at set time intervals and continue to send until either the maximum burst size is reached or a new interval is reached. The intermittent traffic did not use CA, so the usefulness of CA together with uncontrolled traffic could be examined. The CA procedures require data to be sent for a few round-trip times before control of transmit windows is established. Therefore applications that have just started or whose transmissions are of short durations resemble the intermittent traffic used in these experiments. The results indicate that when CA is used by the constant-demand traffic, round-trip times and variability of round-trip times are significantly reduced for all traffic. In these experiments the throughput was also increased somewhat for the constant-demand connections and more dramatically for the intermittent traffic. This indicates that as long as there is a significant component of the network traffic characterized by a constant demand for throughput, CA will produce benefits.

Another set of experiments which was designed to test the generality of CA involved sending two-way traffic through the network configuration. These experiments were originally planned to examine the effect of lost acknowledgements on network performance with and without CA. The unexpected result, however, was that CA reduced fairness in these two-way traffic experiments. Several modifications to the CA procedures aimed at eliminating this lack of fairness were implemented and evaluated. A simple and effective modification was identified by these experiments which provides fairness with two-way traffic that is comparable with the degree of fairness seen with one-way traffic.

The final set of experiments in this report used traffic generators that emulate file transfer and remote login traffic. These experiments were designed to further test the effectiveness of CA in real network environments. As expected from the prior experiments using intermittent traffic, round-trip times, variability of round-trip times, and packet drops were significantly reduced when CA was employed for all traffic. These experiments provide further evidence that CA improves performance for bulk data transfers that require optimal throughput and for interactive traffic that requires low response times.

Experiments with DEC-bit Congestion Avoidance

Rick Wilder, K. K. Ramakrishnan, Allison Mankin

Congestion Avoidance: Implementation and Testing

- Implementation of DEC-developed congestion avoidance algorithms in TP4 and CLNP
  - BSD 4.3 with University of Wisconsin protocols
- Experimentation in the Testbed Network
  - To validate the implementation
  - To demonstrate the effectiveness of the algorithms in a real implementation
  - Try a variety of configurations and traffic types

Summary of DEC-bit Congestion Avoidance (CA)

- Network provides feedback to users indicating resource utilization
  - A single bit in each network header indicates "congestion" or "no congestion"
- Receiving transport protocols sample the bits (over 2 windows) and adjust the credits granted the sender.

Summary of DEC-bit Congestion Avoidance (CA) (continued)

- Router Algorithm
  - Queue lengths are averaged over the previous and current busy periods plus the idle time in between ("busy" means queue length > 0)
  - Congestion bit set in PDUs forwarded while the queue length average > 1
- Transport (User) Algorithm (by RECEIVER)
  - Bits are sampled over two transmit windows
  - If less than half the bits were set, the credits (window size) are increased by one PDU
  - Otherwise credits are reduced multiplicatively (* .0875).

Summary of DEC-bit Congestion Avoidance (CA) (continued)

- Transport Congestion Recovery (CUTE) (by SENDER)
  - When a timeout occurs a congestion window is set to one PDU
  - The congestion window is increased by one after each window of acks are received

Congestion Avoidance Implementation

- Included:
  - Additions to the UNIX kernel-based CLNP and TP4 implementations
  - OSI monitoring software to support experimentation
  - Interface to SLIP driver to support low speed links
  - Modified retransmission behavior to stand up to heavily congested networks
    - based on a single timer per connection
    - uses the Karn-Partridge (clamped backoff) algorithm
  - A "baseline" with which to compare the CA software
    - a fixed window size which is temporarily reduced after a retransmission timeout
    - CUTE (CA in end-systems only)

Conclusions from the Implementation Effort

- CA itself is not difficult to implement
- Setting up testbed environment and doing experiments was very time consuming
- Instruction Counts for Updating Queue Length Statistics

	68020	MicroVAX
Normal Case	15	14
Starting New Cycle	29	26

Experiments Reported

- "Low demand"
  - no. of connections less than optimal window size
- "High demand" (overloaded)
  - no. of connections greater than optimal window size
- Intermittent Traffic
- File Transfer / Remote Login traffic generators
- Two-way Traffic

Initial Testbed Configuration

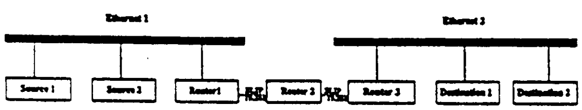

Low-demand Experiments

- CA algorithms behaved as anticipated
  - Window sizes responded to different load levels to fairly distribute throughput
  - Router utilization was high while queue lengths remained small

Window Size for a Single Connection

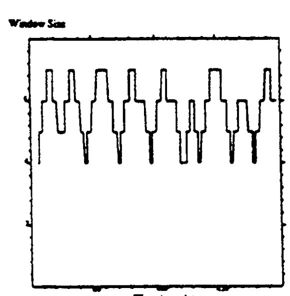

Window Sizes for Two Connections

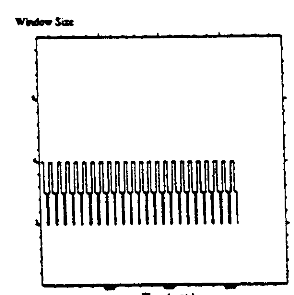

Data Transmissions for Four Connections

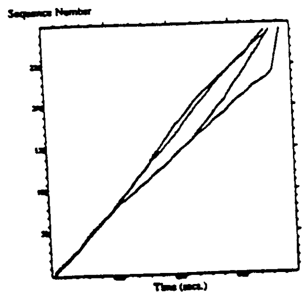

High-demand Experiments

- CA reduced window sizes to avoid retransmissions and keep router queues short
- Dramatic results in terms of fairness and response time was seen
- Throughput was better for CA but the differences were small with this experimental configuration

Data Transmission for Ten Connections: Baseline

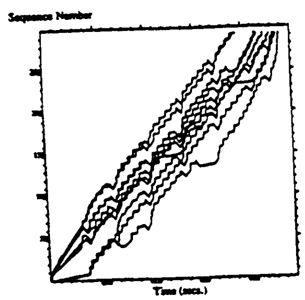

Data Transmission for Ten Connections: CA

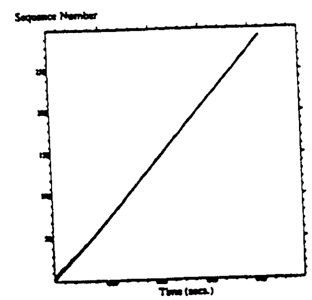

Round Trip Delays for Two Representative Connections

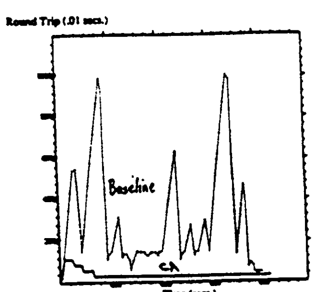

Round Trip Delay and Retransmissions for 10 Connections

	From First Trans.			From Last Trans.	
	Mean	Std. Dev.	Retrans.	Mean	Std. Dev
CA:	3.69	2.15	0	3.69	2.15
Baseline:	19.86	8.97	73	19.54	7.74

Intermittent Traffic Experiments

- Connections with a constant throughput demand share a router with transient-demand connections from another source network
  - transient-demand connections never used CA
  - CA and "baseline" algorithms were used with constant-demand connections
- CA expected to control queue lengths to minimize drops
- Topology allows us to load more of the routers
- Tests the ability of CA to handle uncontrolled bursts of traffic
  - New connections starting up
  - Connections with a short or fluctuating demand

Intermittent Traffic Experiments (continued)

- Two traffic mixes were used
  - 6 constant-demand plus 2 transient-demand
  - 2 constant-demand plus 2 transient-demand
- A large reduction in drops and delay was seen for both traffic mixes
  - Benefits of CA are not limited to:
    - highly congested networks
    - networks where a very high percentage of the traffic is controlled (constant-demand with CA)

Network Configuration for Intermittent Traffic Exp.

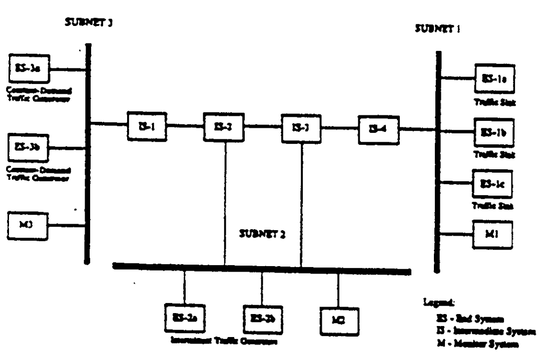

Data Transmissions for 6 Constant + 2 Intermittent, Baseline

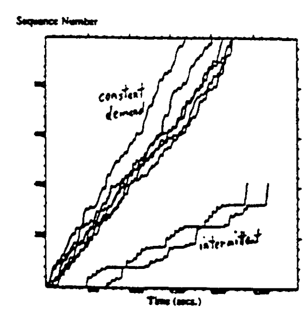

Data Transmissions for 6 Constant + 2 Intermittent, CA

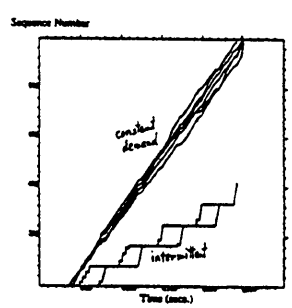

Round-Trip Delays and Retransmissions (6 + 2)

	From First Trans. Mean	From First Trans. Std. Dev.	Retrans.	From Last Trans. Mean	From Last Trans. Std. Dev
CA					
Intermittent:	4.84	4.73	17	4.59	3.94
Constant:	3.53	1.94	13	3.50	1.71
BASELINE					
Intermittent:	14.13	19.84	179	6.72	10.50
Constant:	10.70	10.38	385	9.47	7.73

Data Transmissions for 2 Constant + 2 Intermittent, Baseline

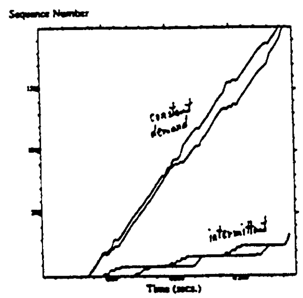

Data Transmissions for 2 Constant + 2 Intermittent, CA

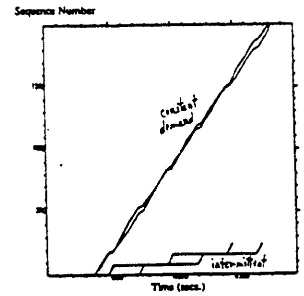

Round-Trip Delays and Retransmissions (2 + 2)

	From First Trans.		Retrans	From Last Trans.	
CA	Mean	Std. Dev.		Mean	Std. Dev
Intermittent:	3.87	.78	0	3.87	.78
Constant:	2.51	.98	2	2.50	.86
BASELINE					
Intermittent:	7.46	8.85	40	6.27	6.86
Constant:	5.35	2.78	28	5.31	2.55

Remote Login and File Transfer Experiments

- Intended to increase our confidence in CA's ability to handle "real world" traffic patterns
- Traffic generators emulate remote login and file transfer sessions
  - We expected CA to minimize delay especially for the remote login traffic

File Transfer / Remote Login Traffic Generators

- all emulations run for 2000 seconds
- File Transfer:
  - "File" size, 40K to 100K bytes
  - "Think" time: 20 - 100 seconds
- Terminal Traffic:
  - PDU size: 20 - 80 bytes
  - "Think" + response time: 10 - 20 seconds
- Host Traffic:
  - PDU size: 20 - 2000 bytes
  - "Think" + response time: 10 - 20 seconds

Data Transmissions for Remote Login / File Transfer, Baseline

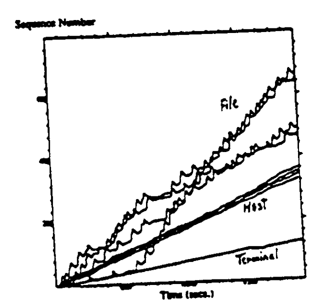

Data Transmissions for Remote Login / File Transfer, CA

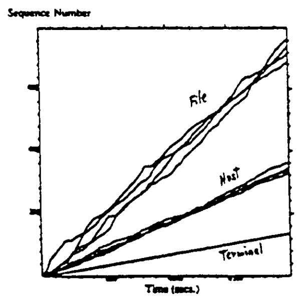

Round-Trip Delays and Retransmissions, File Transfer / Remote Login

CA	From First Trans.		Retrans.	From Last Trans.	
	Mean	Std. Dev.		Mean	Std. Dev
Host:	4.53	1.61	0	4.53	1.61
Terminal:	5.05	1.93	1	5.05	1.93
File:	5.52	2.58	8	5.45	2.06
BASELINE					
Host:	6.37	3.85	5	6.30	3.49
Terminal:	7.56	11.32	58	7.16	10.98
File:	34.34	40.51	729	11.76	19.17
CUTE					
Host:	7.00	2.65	3	7.00	2.65
Terminal:	7.10	4.10	20	7.10	4.10
File:	12.05	13.15	364	8.54	6.61

File Transfer Throughput

- PDUs Sent by 4 connections in 2000 seconds
  - 2275 using Baseline
  - 2660 using CUTE
  - 2920 using CA
- 28% throughput advantage for CA over Baseline

Two-way Traffic

- Each end system will be both a source and a destination for data
- Both data PDUs and ACKs will be discarded during severe congestion
- Drops may affect overall throughput much more than with on-way experiments
  - Drops may cause retransmissions over slow links rather than over ethernet

Two-way Traffic (continued)

- The results showed an unexpected degree of unfairness with two-way traffic
  - The unfairness persisted using different numbers of connections in each direction
    - 2 connections vs. 1 connection was taken as the simplest case for analysis
  - Unfairness turned out to be due to ACKs from the reverse direction creating "small" congestion cycles
    - short cycles cause long cycles to be quickly forgotten
  - The data windows for competing connections tend to be sent in bursts
  - The ordering of connections in bursts tends to persist
    - Those at the beginning of the bursts "win"
    - Those at the end of the bursts "lose"

2 Connections vs. 1, Unmodified CA

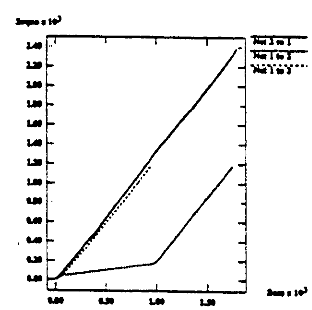

Queue Events for Unmodified CA in 2 Connections vs. 1

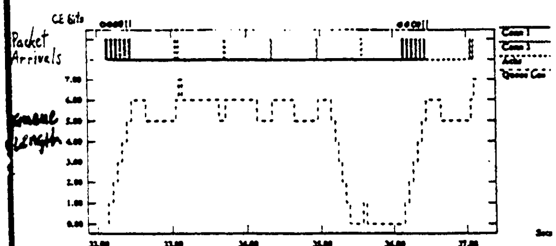

Modification to CA to improve fairness with two-way traffic

- Record the number of PDUs processed during each congestion cycle
- At the end of a congestion cycle, check the number of PDUs processed
  - if PDU count is one, combine this cycle with the "previous cycle" (rather than discarding the "previous cycle")
    - add the area under the curve for the 2 cycles
    - make the end time for the combined cycle be the end time for the "short cycle"

Queue Events for Modified CA in 2 Connections vs. 1

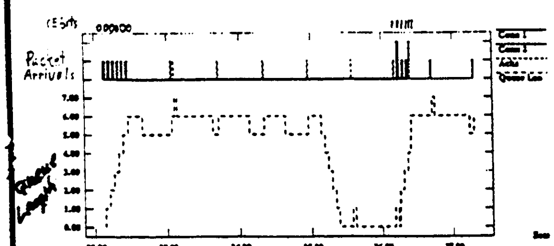

2 Connections vs. 1, Modified CA

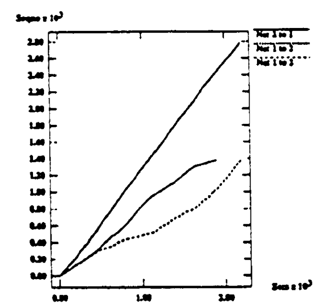

4 Connections vs. 4, Modified CA

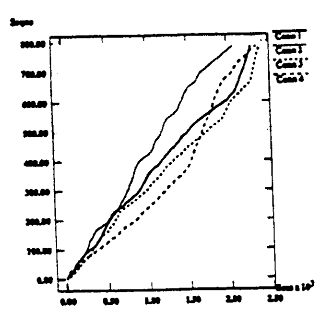

4 Connections vs. 4, Modified CA

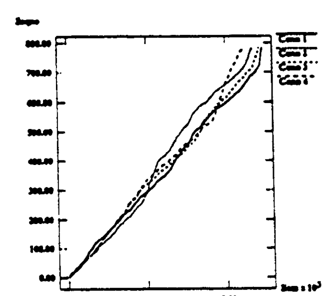

Performance of Modified CA with Two-way Traffic

	Delay		Retrans.	Time to Complete 3.2MB
	mean	Std Dev.		
CA (Modified)	6.10	1.64		
	6.13	1.78	11	2467 sec.
CUTE*	14.62	13.25		
	14.92	13.75	864	2540 sec.
Baseline	16.26	19.68		
	18.06	22.04	750	2610 sec.

Conclusions from Two-Way Traffic Experiments

- May not be a serious real-world problem
  - File transfer / remote login experiment did not exhibit a serious problem
  - Two-way traffic plus intermittent sources did not show a problem

Further Work Required

- Experiments using higher link speeds
  - Interest on the part of DARPA to do this in DARTNET
- Different PDU and window sizes
- Standards
  - Communicate modification of CA for two-way traffic to standards groups
  - See if CA can be made mandatory for NIST Implementation Agreements and/or GOSIP
    - Use of QOS option for CLNP and use of CA algorithms are separate issues

6.4 Proposed Internet Routing Protocol Standardization Criteria

Presented by Bob Hinden/BBN

PROPOSED INTERNET ROUTING PROTOCOL STANDARDIZATION CRITERIA

Bob Hinden

March 12, 1990

BBN Communications

March 12, 1991

PRESENTATION OVERVIEW

- Motivation
- General Requirements
- Requirements for
  - Proposed Standard
  - Draft Standard
  - Full Standard
- Issues

BBN Communications

March 12, 1991

MOTIVATION

- Routing Protocols are More Complex than Most Other Protocols
  - Widely Distributed Real Time Algorithms
  - Implementation and Testing Difficult
- Reduce the Risk that there will be Serious Technical Problems with New Routing Protocols
- Insure that new Routing Protocols will support the continued Growth of the Internet

BBN Communications

March 12, 1991

GENERAL REQUIREMENTS

- Specifications Describing Protocol and Usage
  - Well Written and Understandable
- Management Information Base
- Security Architecture
  - Authentication is Key Issue
- Interoperable Implementations
- Testing of all Features of the Protocol
- Operational Experience
- Report Documenting Testing Results and Operational Experience
- Analysis of the Routing Protocol to show how it will Support Internet Growth

BBN Communications

March 12, 1991

REQUIREMENTS FOR PROPOSED STANDARD

- Well Written Protocol Specification and Usage Documents
- MIB for Routing Protocol in Internet Draft Form
- Security Architecture for the Protocol Specified
- One or More Implementations
- Testing of Major Features of the New Routing Protocol

BBN Communications

March 12, 1991

REQUIREMENTS FOR DRAFT STANDARD

- Revised Protocol and Usage Documents based on Additional Experience
- MIB at the Proposed Standard Level
- Two or More Interoperable Implementations
  - Two Must have been Written Independently
- Testing of all Features of the Protocol
  - Including Security Features
- Significant Operational Experience in the Operational Internet
  - Moderate Number of Routers
  - Moderate Complex Topology
- Report Documenting Testing Results and Operational Experience
- Report Analyzing the Routing Protocol
  - Resources Consumed
  - Scaling as Routing Environment Grows
  - Limits of Protocol

BBN Communications

March 12, 1991

REQUIREMENTS FOR FULL STANDARD

- Revised Protocol and Usage Documents based on Additional Experience
- MIB at the Full Standard Level
- Three or More Interoperable Implementations
  - Two Must have been Written Independently
- Testing of all Features of the Protocol
  - Between Independent Implementations
- Significant Operational Experience in the Operational Internet
  - Large Number of Routers
  - Complex Topology
  - Multi-Vendor Operation
- Report Documenting Testing Results and Operational Experience
- Update to Report Analyzing the Routing Protocol

BBN Communications

March 12, 1991

ISSUES

- Too Hard or Too Easy?
- Should Draft Standard or Full Standard be Biggest Hurdle?
- How to Specify Operational Experience
  - Does Moderate Equal 5, 10, 100 Routers?
- Form of Reports?
  - RFC's, Presentations?
- Level of Detail in Reports?
- Should Simulations be Required?

BBN Communications

6.5 BIND 4.8.4...and Beyond

Presented by Philip Almquist/Consultant

BIND is by far the most widely used implementation of the domain name service (DNS). Despite its widespread usage, BIND has limitations and bugs that are all too well known to many system and network managers. Therefore, I am pleased to announce that I am beginning a major effort to improve BIND.

I have begun by developing a tentative project plan. This talk is a brief overview of that plan, presented in order to encourage feedback from any of you who are interested. That feedback will help me to determine where revisions in the tentative plan are needed.

The second of the accompanying slides enumerates the goals of the project. Subsequent slides list the most important features planned for each software release. The final slide is the tentative schedule.

I hope that other people will contribute additional improvements. I will do the necessary work to integrate useful contributions into the released software.

Beginning with BIND 5.0, I will be supplying a low level of software support. I will attempt to provide patches promptly for problems which are reported using a form that can be submitted using electronic mail. Patches will available via anonymous FTP, and maintenance releases will be created when necessary.

Thanks to Dave Crocker, this project is being supported as a public service by DEC's Network Systems Laboratory.

BIND 4.8.4...and Beyond

Philip Almquist

almquist@Jessica.Stanford.EDU

214 Cole Street, Suite #2, San Francisco, CA 94117-1916

I am grateful to Digital Equipment Corporation's Network Systems Lab for providing support for this project.
I am neither an employee of nor a spokesperson for DEC.

Outline of this Talk

1. Introduction
2. Project goals
3. Project milestones
4. Tentative schedule

Goals (Why Fix BIND)

- Make it conform to relevant standards
- Decrease unnecessary Internet congestion
- Make it easier to operate and manage
- Don't require an accessible root server
- Don't let *named* originate, believe, or propagate bogus information
- Add hooks for administrative policy controls
- Make the code clean and robust
- Provide software support
- Encourage and incorporate work by others

BIND 4.8.4

Interim, unsupported release:

- Changes requested by DNS WG to reduce bogus data
- Changes to record the source of cached data
- Access list for zone transfers
- A few bug fixes

BIND 5.0

Supported major release:

- Extensive internal changes
- Negative cacheing
- RFC-1183 RR type support
- Fix for "European Root Server Problem"
- Several less important changes

BIND 6.0

Supported major release:

- Standards-conformant *named*
- User-defined access controls
- No requirement for accessible root servers
- Ease of operation improvements
- Erroneous data suppression
- Several less important changes

BIND 7.0

Supported major release:

- Robust, standards-conformant resolver
- Better duplicate suppression
- SNMP support
- Other goals TBD

Tentative Schedule

- 4/91: BIND 4.8.4
- 2/92: BIND 5.0
- 10/92: BIND 6.0
- Q2/93: BIND 7.0
- Q2/94: final maintenance release

6.6 NREN Legislative Update

Presented by William Bostwick/DOE

The HPCC Program

William Bostwick

Federal Networking Council (FNC)

April 4, 1991

HPCC Program

- Presidents FY1992 Budget
- OSTP
  - FCCSET
  - Committee on Physics, Math and Engineering Science (PMES)
- Defines Responsibilities
- Management and Reporting

Discussion on Authorization

- Why Authorizations?
- Can be helpful/Can constrain
- Subcommittee Positions

Authorization Hearings

- Senate Commerce Committee - S272
  - March 5th, 1991
- House Science and Technology - HR-656
  - March 19, 1991
- Senate Energy Committee - S343
  - April 11, 1991

What's Important

- Appropriations
- Unity
  - Feds
  - Community

To Obtain Copy of "Blue Book"

Write:
PMES
c/o NSF
Computer and Information Science and Engineering
1800 G Street, NW
Washington, DC 20550

Or call:
NSF (202) 357-7936
DoE (301) 353-5800
DARPA (703) 614-5800

6.7 Advanced Networks Research at Washington University

Presented by Jonathan Turner/WashU

The advanced networks research program at Washington University includes a blend of experimental, analytical and theoretical work with a strong systems focus. We are committed to the proposition that successful engineering research requires an intimate knowledge of the practical issues involved in building complex systems, as well as strong analytical capabilities. We find that theoretical research, in the absence of a strong connection to practical concerns, too easily drifts into activities of interest only to an ingrown research community, and is ultimately sterile and unprofitable. On the other hand, experimental work that is uninformed by a deeper understanding of fundamental issues, can have only limited and short-term value.

Over the past several years, our research activities have centered on packet switching systems operating at link speeds of about 100 Mb/s. This work has included the design and implementation of integrated circuit components and their use in experimental switching systems; performance evaluation of switching systems from a variety of different points of view; design of connection management protocols for multicast networks; high speed internetworking and host-network interfaces; televisualization over wide-area high speed networks; performance evaluation of broadband networks; design of multicast routing protocols; buffer and bandwidth management; distributed debugging systems; special purpose computer-aided design tools; and video coding algorithms.

In this talk, I will summarize some of our work on switching system architecture, describe a four node demonstration network being constructed in the St. Louis metropolitan area to establish the technical feasibility of our approach and explore its use in visual communication applications. I will also describe some of our plans for Project Zeus, a proposed application of fast packet switching technology to a general purpose campus network.

Advanced Networks Research

Jon Turner
Computer Science Department
Washington University

This work supported by:
- National Science Foundation (grant DCI 8600947)
- Bell Communications Research
- Bell Northern Research
- Italtel SIT
- NEC
- NTT
- Southwestern Bell

Advanced Networking Research at Washington University

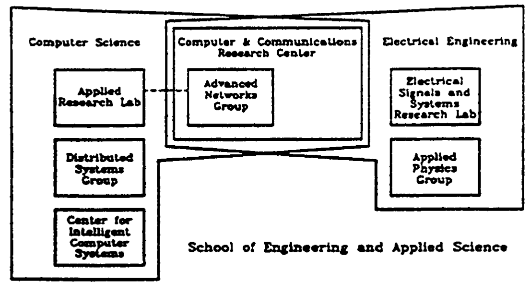

Advanced Networks Group operates within the Computer and Communications Research Center, which is in overlap of CS and EE departments; 3 faculty, 13 graduate students and 1 visitor.

ANG is concerned with the design and analysis of flexible, high performance networks, with emphasis on multicast communication.

Applied Research Lab (ARL) is concerned with transfer of faculty research results to industrial practice; full-time staff of 7 plus 3 visitors.

Networking Research Topics

Switching System Design and Analysis (Turner)
: Haifeng Bi, Rick Bubenik, Lorenzo Favalli Riccardo Melen, Einir Valdimarsson, Ellen Witte

Demonstration Network Design and Construction (Turner)
: Neil Barrett, Pierre Costa, Gaurav Garg, Rex Hill, Shabbir Khakoo, Noritaka Matsuura, Nader Mirfakhraei, George Robbert, Randy Richards, Kenichi Sato, James Sterbenz, Einir Valdimarsson

Multicast connection management (Turner)
: Rick Bubenik, John DeHart, Mike Gaddis, Victor Griswold, Kurt Haserodt, Mark Hunter, Ken Katsumata, Dave Wexelblat

Connection-Oriented Internetworking (Parulkar)
: Millind Buddhikot, Zubin Dittia, Chuck Cranor, Tony Mazraani

High Speed Host-Network Interface (Parulkar)
: Anshul Kantawala, James Sterbenz

Visualization on High Speed Networks (Parulkar)
: Larry Gong, Christos Papadopoulos

Performance Evaluation of Broadband Networks (Bovopoulos)
: Diamantis Kotoulas, Einir Valdimarsson

Resource Management (Bovopoulos, Turner)
: Shahid Akhtar, Akira Arutaki

Diversity and Change Require Flexible Networks

Current networks tailored to specific applications; difficult to adapt.

Flexibility needed to adapt to

- diverse and changing application mix
- diverse and changing technology base

Visual applications (image/video), multimedia workstations driving need for greater speed.

Objectives for advanced networks

- application-independent transport
- channels of arbitrary bandwidth (1 b/s and up)
- high bandwidth access and trunk facilities (.1–10 Gb/s).
- *general* multipoint connections
- inexpensive and ubiquitous connections.

Multipoint Connections

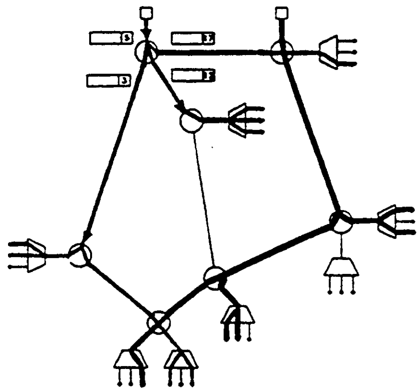

Broadcast trees grow and shrink dynamically.

Bandwidth of source and number of receivers completely flexible.

Conference connection—multi-way broadcast.

Switch Architecture

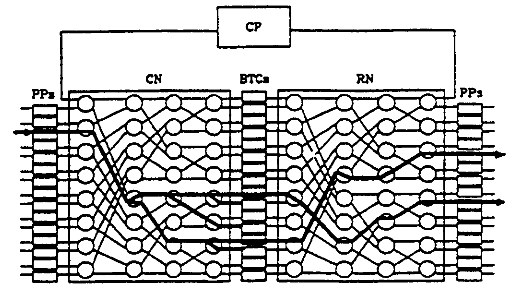

Routing Network (RN)—routes packets to proper output.

Copy Network (CN)—replicates broadcast packets.

Broadcast Translation Circuits (BTC)—add routing information to each copy of broadcast packets.

Connection Processor (CP) controls operation of system.

In prototype, internal data paths are eight bits wide, 25 Mbs clock rate.

Packet Processor

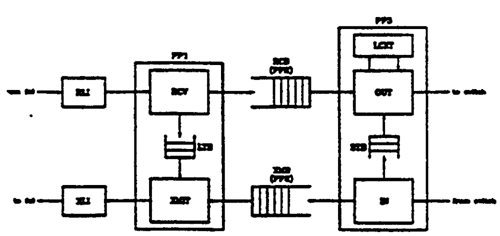

Buffers: Receive Buffer, Transmit Buffer, Link Test Buffer and Switch Test Buffer.

Logical Channel Translation Table (LCXT) is used to determine routing of each packet; maintained by CP.

Output circuit performs logical channel translation, adding extra header to incoming packets.

Prototype implementation requires four chips.

Copy Network and Broadcast Translation

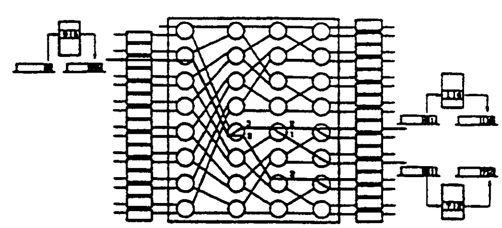

Broadcast packets contain *Fanout* field and *Broadcast Channel Number* (BCN).

Packets replicated based on Fanout value—late splitting.

BCN used to index table in BTCs. Each copy routed to distinct output.

Demonstration Network

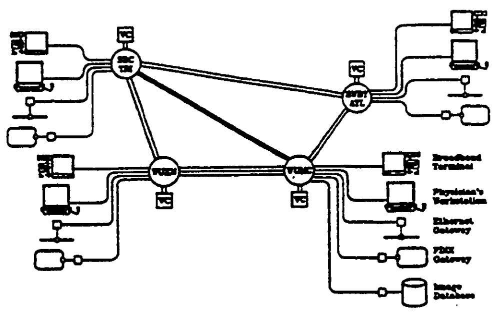

Application of fast packet switching to visual communication.

Four site demonstration network in St. Louis.

Image and video compression.

Physician's workstation and associated medical applications.

Compatible with emerging ATM standards.

Linkage to similar experiments anticipated.

Sponsored by Southwestern Bell and NEC America.

Medical Image Network

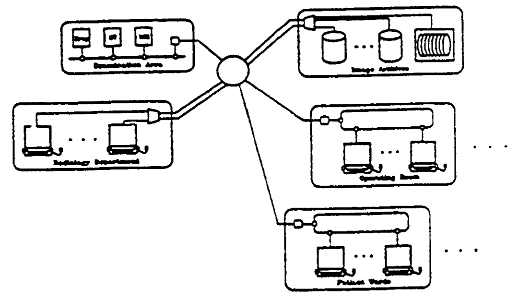

Large medical center with several hosptials, medical school, associated clinics and research facilities.

Generate several thousand diagnostic images daily.

Variety of image types

digital xrays	nuclear medicine
CT, MRI studies	volume rendering

Variety of viewing situations

radiologists; high volume
ward and operating rooms
remote clinics

Prototype System Organization

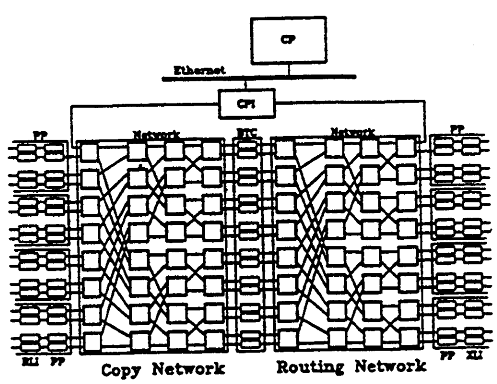

Copy and routing networks are implemented with single board having 32 PSE chips.

Broadcast translation circuit board includes all 16 BTCs.

Packet processors boards include 4 PPs each, including optical transducers. packets to outgoing PPs.

Connection processor interface board prcvides

Network Board

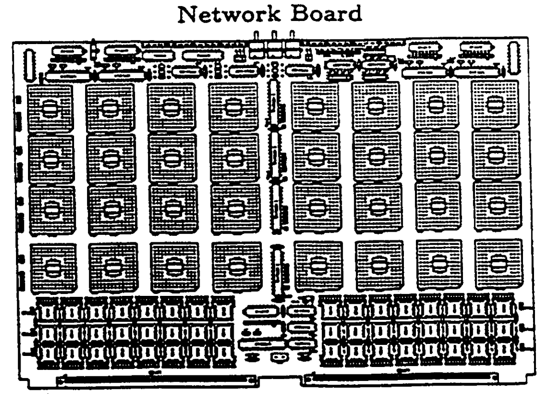

Contains 16 port network. PSE chips packaged in 108 pin PGAs.

High density connectors provide almost 800 pins at board level; 320 data signals plus associated grounds consume 640.

Clocked latches provide high speed operation across backplane.

Timing circuitry generates local timing signals to synchronize board with global timing cycle.

Ethernet Portal

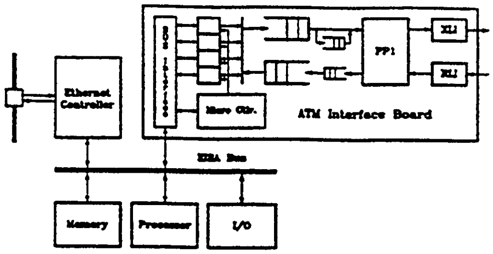

ATM Interface Board will be EISA Bus compatible and plug into standard EISA PC.

Ethernet Controller transfers packets to/from dual-ported memory on ATM Interface Board.

Microcontroller transfers packets between memory and fifos.

Main processor manages memory and Ethernet controller's buffer list to accomplish fragmentation and reassembly.

Initial firmware implementation will transfer packets to/from static multicast connection; filtering and use of multiple connections can be added later.

Expected completion date: 2Q91.

Broadband Terminal

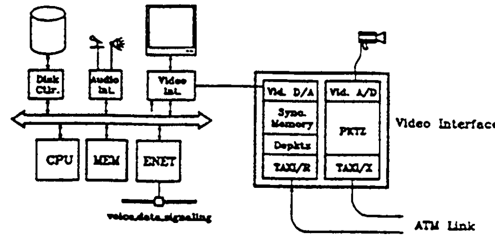

Digital video carried uncoded over ATM link.

Voice, data and signaling carried to ATM network through Ethernet portal.

Audio and video interfaces bypass bus.

Expected completion date: 2Q91.

Project Zeus
Design of a Fast Packet Network
and its Application on a University Campus

Jon Turner
Applied Research Laboratory
Washington University, St. Louis

Motivation and Objectives

Motivation

- ATM technology has been accepted for next generation networks.
- Demand for high speed applications will come from computers and multimedia workstations.
- LAN and workstation vendors recognizing need for switched campus networks.
- Washington University offers unique opportunities for creation and application of fast packet campus network.

Objectives

- Stimulate commercial development of network components needed for a fast packet campus network supporting thousands of workstations.
  - Design, develop and apply prototype hardware and software.
  - Provide context in which multiple vendors can cooperate to create new technology and develop new markets.
  - Help develop consensus needed to accelerate standards.
- Explore applications of fast packet campus nets through establishment of a network supporting range of applications in research, clinical medicine and education.

A Fast Packet Campus Network

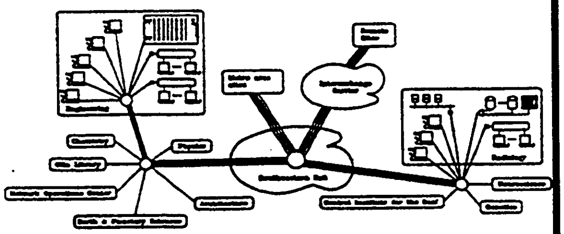

Mixed network with shared access LANs and directly connected workstations coupled to ATM core.

2-5 ATM switches with 64-256 ports and richly interconnected.

Dozens of shared access subnets and hundreds of multimedia workstations.

Access to NSF-net backbone at 45 Mb/s for scientific applications.

Local companies accessed through public network over DS-3 or Sonet; preferably switched ATM over OC-3.

Clinical office building at medical center planned to open in 1992-3; image communication as utility.

Project Zeus Timetable

Phase 0 — Current Fast Packet Project (1989–1991)

- Demonstrate feasibility of core technology.
- Provide basis for more complete design and implementation.
- Provide testbed for application development.

Phase 1 (1991–1993)

- Initiate application development using phase 0 network.
- Develop essential components needed to establish core campus network and support application development
- Establish basic interoperability between campus and public network ATM switches.
- Initiate public network trial of medical imaging.

Phase 2 (1993–1995)

- Improve economy of implementation; develop components for larger scale, higher speed networks.
- Improved interfaces to existing LANs.
- Complete interoperability between campus and public networks.
- Expand public network trial of medical imaging.
- Deploy range of applications throughout campus.

Phase 1 Benchmark Network

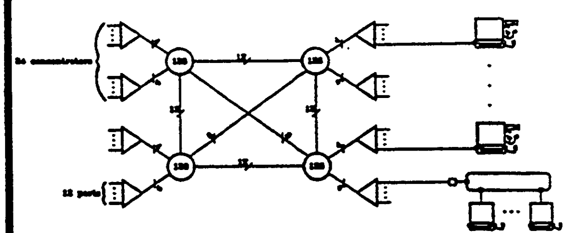

Supports over 1000 ports at 150 Mb/s each.

Interfaces for multimedia workstations, Ethernet and FDDI.

Inter-switch and switch-to- concentrator links use 4 × 150 Mb/s interface.

Control Processors in each switch control concentrator configuration.

Resource allocation based on peak bandwidth requirements.

Phase 1 Switching System

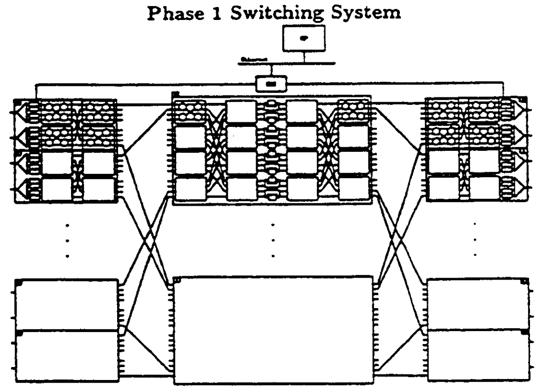

Packet Processor board includes 8 PPs with link interfaces and two eight port switch sections; four links multiplexed on single fiber.

Network board includes 16 port copy network section, 16 port routing network section and Broadcast Translation Circuits.

Binary Packet Switch Elements packaged four per chip.

PP requires one custom chip plus three memory chips.

BTC controllers, 2 per chip; 2 BTCs share memory chip.

Each link supports 4096 individually switchable VPIs and 4096 VCIs; system supports total of 4096 multicast connections.

16 to 128 ports; n port configuration requires $1 + 3n/16$ cards.

Packet Processor

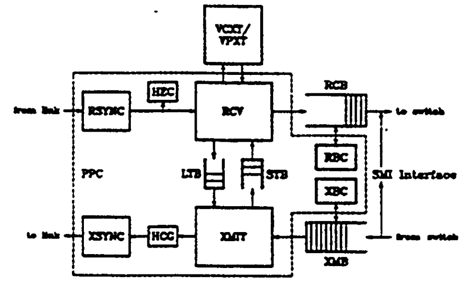

PP Controller (PPC) is custom chip with estimated 150K transistors.

VCXT/VPXT is 9 × 32K SRAM giving 4K VPI entries and 4K VCI entries; packets with VPI=0 are routed on VCI, remainder on VCI.

RCB is 9 × 4K fifo; RBC includes packet counter.

XMB is pair of 9 × 8K SRAMs; XBC includes resequencing logic.

Synchronizers allow switch to support any mix of link speeds up to 155 Mb/s.

6.8 Electronic Radiology Laboratory

Presented by Dr. G. James Blaine/Washington University Medical School

Tour of

Electronic Radiology Laboratory

Dr. G. James Blaine

Associate Professor of Computing Sciences
Department of Radiology
Co-director Electronic Radiology Laboratory
Mallinckrodt Institute of Radiology
Washington University Medical School

ELECTRONIC RADIOLOGY LABORATORY

Mallinckrodt Institute of Radiology
Washington University Medical Center
St. Louis, Missouri

The MIR has established a Radiology Image and Information Management testbed program centered within the Electronic Radiology Laboratory (ERL). The program encompasses research on image compression, soft-copy display, as well as the performance measurement and modeling of picture archiving and communication systems. The testbed program is supported by an environment rich in new subsystems that are being evaluated in collaboration with major suppliers of radiology equipment. These include high performance displays, a high performance film scanner/printer, an advanced multiviewer, specialized storage systems, and low cost electronic viewboxes. Collaborators include Digital Equipment Corporation, Siemens Medical Systems, Philips Medical Systems, Eastman Kodak Company, Southwestern Bell Corporation - Technology Resources, Inc. and Imlogix, Inc.

The testbed is a distributed environment for testing and development where coupling to clinical activities is easily varied. Extensive Ethernet-based network facilities provide access to the clinical acquisition instruments. Software is being developed to support routine image acquisition, storage in an on-line archive, and effective retrieval at a Physician's Inquiry & Display Station. This LAN-based system utilizes DEC equipment for the image database system. The software environment is currently based on VMS and utilizes the DECimage Applications Services, DECwindows, DECnet and Rdb software packages. An extensive software development, MIR-RIM, supports image acquisition, access to the image database as well as our operational Radiology Information System. Teleradiology is the subject of our collaboration with SBC Technology Resources, Inc.

The tour will include a brief program overview followed by demonstrations of a high-performance film scanner, an Inquiry & Display Station for primary interpretation, and a teleradiology Inquiry & Display Station using narrow-band ISDN.

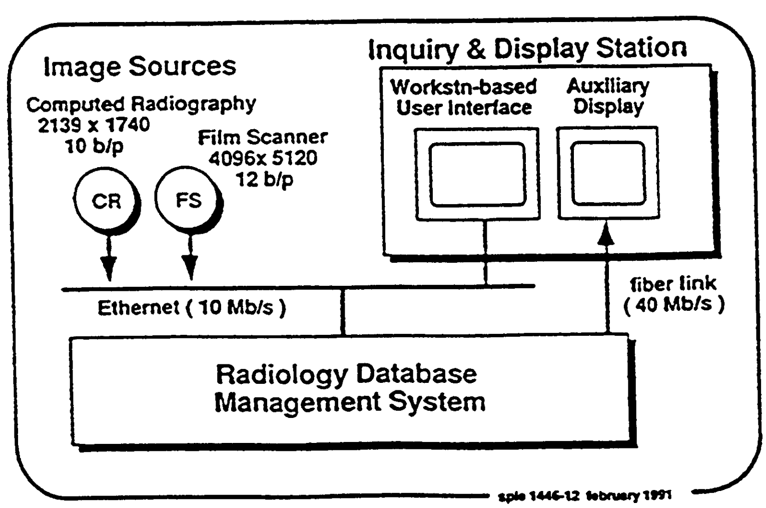

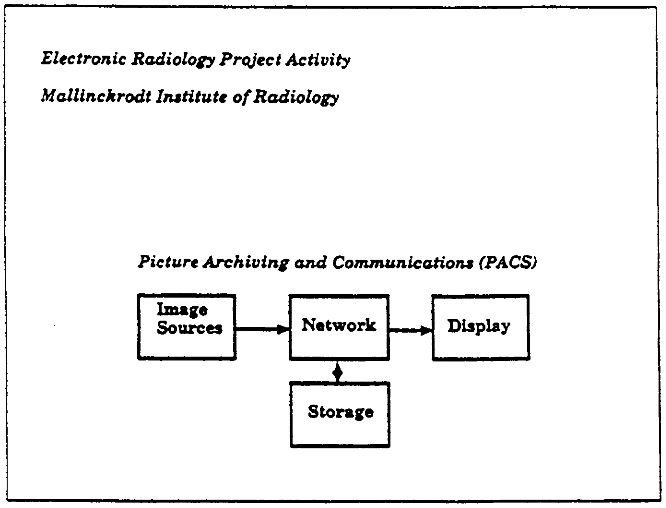
Electronic Radiology Project Activity
Mallinckrodt Institute of Radiology
Picture Archiving and Communications (PACS)
Image Sources
Network
Display
Storage

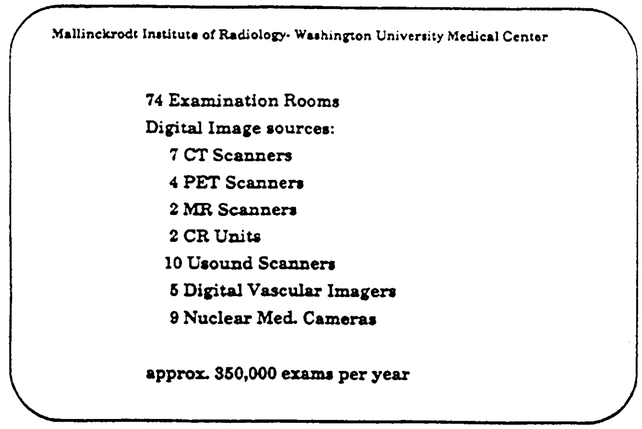
Mallinckrodt Institute of Radiology- Washington University Medical Center
74 Examination Rooms
Digital Image sources:
7 CT Scanners
4 PET Scanners
2 MR Scanners
2 CR Units
10 Usound Scanners
5 Digital Vascular Imagers
9 Nuclear Med. Cameras
approx. 350,000 exams per year

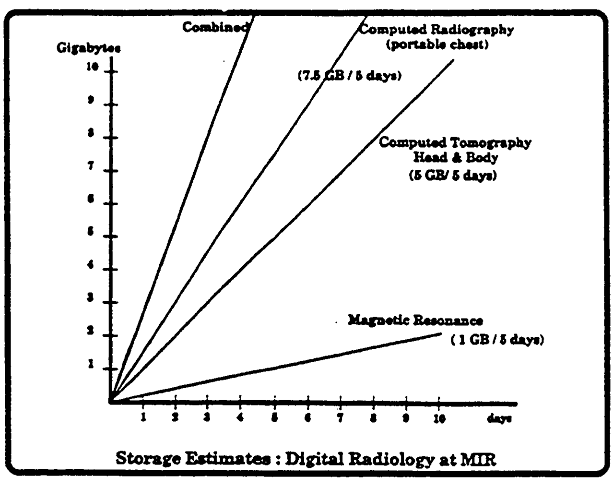
Gigabytes
Combined
Computed Radiography (portable chest)
(7.5 GB / 5 days)
Computed Tomography Head & Body (5 GB/ 5 days)
Magnetic Resonance (1 GB / 5 days)
days
Storage Estimates : Digital Radiology at MIR

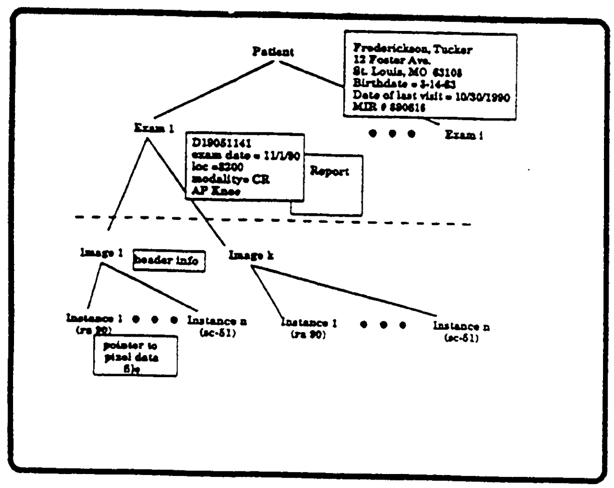
Patient
Exam 1
Exam i
Report
Image 1
header info
Image k
Instance 1
Instance n
pointer to pixel data

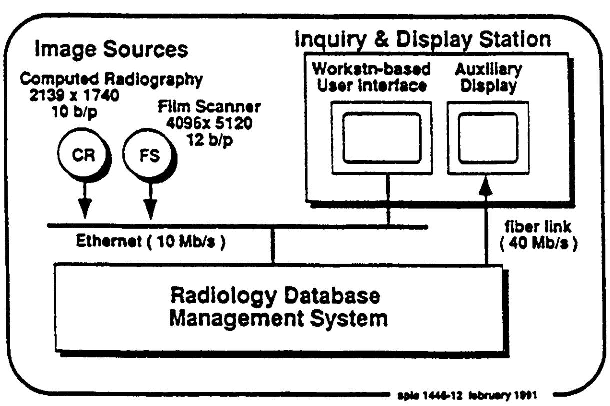
Image Sources
Computed Radiography
2139 x 1740
10 b/p
Film Scanner
4096x 5120
12 b/p
CR
FS
Inquiry & Display Station
Workstn-based User Interface
Auxiliary Display
Ethernet (10 Mb/s)
fiber link (40 Mb/s)
Radiology Database Management System
spie 1446-12 february 1991

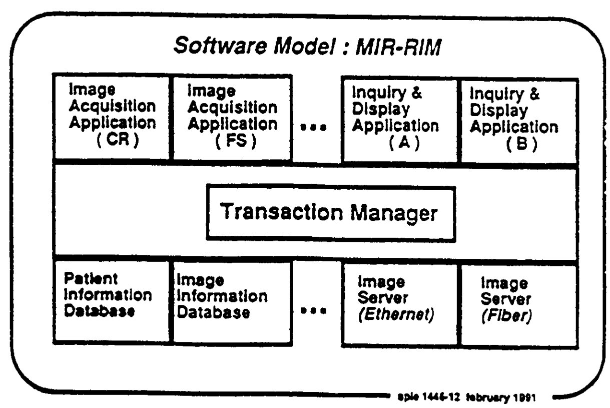
Software Model : MIR-RIM
Image Acquisition Application (CR)
Image Acquisition Application (FS)
Inquiry & Display Application (A)
Inquiry & Display Application (B)
Transaction Manager
Patient Information Database
Image Information Database
Image Server (Ethernet)
Image Server (Fiber)
spie 1446-12 february 1991

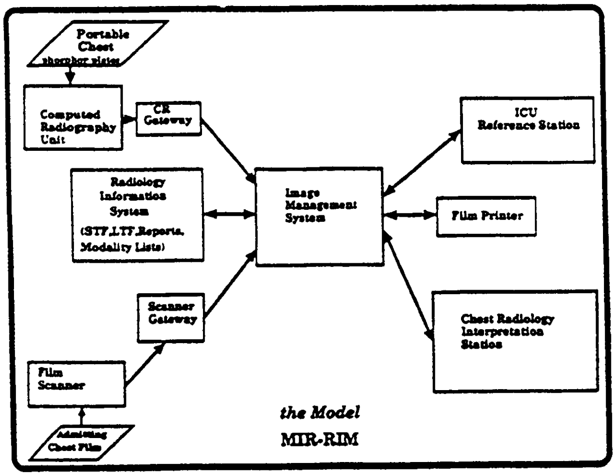
Portable
Chest
Computed
Radiography
Unit
CR
Gateway
Radiology
Information
System
(STF,LTF,Reports,
Modality Lists)
Image
Management
System
ICU
Reference Station
Film Printer
Chest Radiology
Interpretation
Station
Scanner
Gateway
Film
Scanner
Admitting
Chest Film
the Model
MIR-RIM

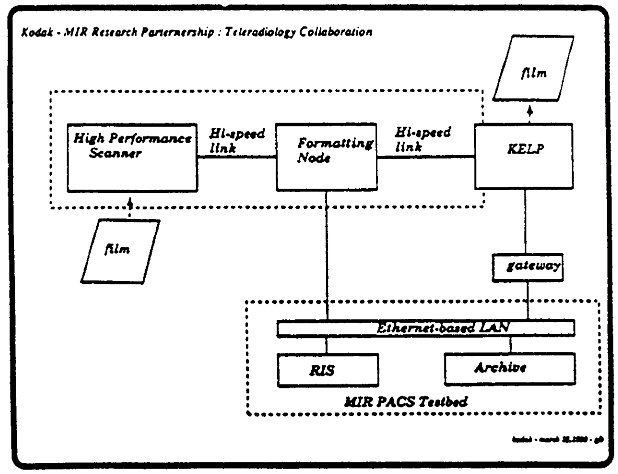
Kodak - MIR Research Partnership : Teleradiology Collaboration
film
High Performance
Scanner
Hi-speed
link
Formatting
Node
Hi-speed
link
KELP
film
gateway
Ethernet-based LAN
RIS
Archive
MIR PACS Testbed

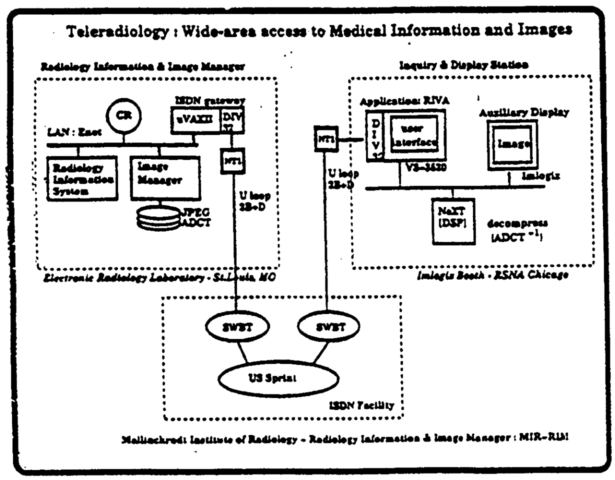
Teleradiology : Wide-area access to Medical Information and Images
Radiology Information & Image Manager
Inquiry & Display Station
ISDN gateway
CR
uVAXII
LAN : Enet
Radiology
Information
System
Image
Manager
JPEG
ADCT
NT1
U loop
2B+D
Application: RIVA
Auxiliary Display
user
interface
Image
VS-3520
Imlogix
NeXT
(DSP)
decompress
Electronic Radiology Laboratory - St.Louis, MO
Imlogix Booth - RSNA Chicago
SWBT
SWBT
US Sprint
ISDN Facility
Mallinckrodt Institute of Radiology - Radiology Information & Image Manager : MIR-RIM

6.9 IETF Access to SMDS

Presented by Kirk Williams/Southwestern Bell Technology Resources

SOUTHWESTERN BELL TECHNOLOGY RESOURCES

Kirk Williams
Senior Technologist
Advanced Networking Technology
550 Maryville Center
St. Louis, MO 63141
314 529-7586
kirk@sbctri.sbc.com

SMDS Access to IETF Acknowledgements

Southwestern Bell Technology Resources
- Greg Newell
- Kirk Williams

Southwestern Bell Telephone Company
- Steve Crider
- Marty Tanner

Washington University
- Martin Dubetz
- Chris Myers
- Guru Parulkar

IETF Internet Access via SMDS

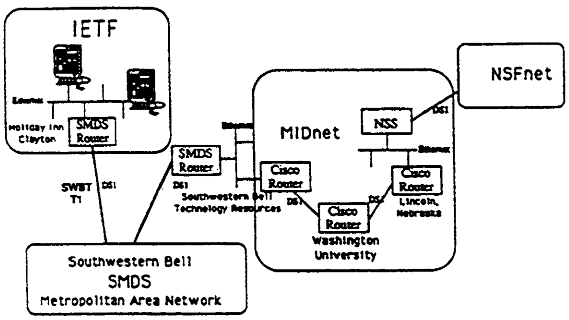

IETF Internet Access via SWBT SMDS Testbed

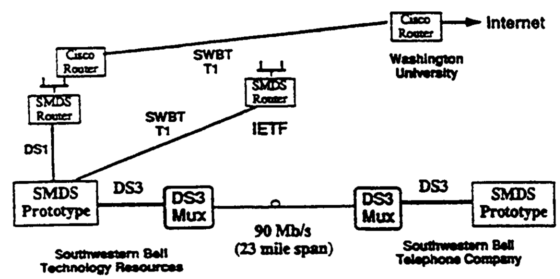

SMDS Broadband Service: INTEROP '90 Demonstration Network

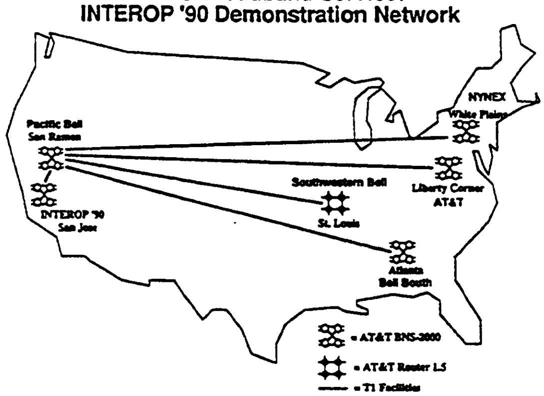

SMDS Interest Group

For information contact:
Elaine Kearney @ 415 949-1779

Bellcore Document Center
908 699-5800
201 292-0067 fax

6.10 Introduction to Frame Relay

Presented by Andy Malis/BBN

This talk provides a basic introduction to the subject of Frame Relay. It discusses the marketing forces and technology trends that resulted in Frame Relay, it describes Frame Relay and related standards activities, it provides examples of how Frame Relay can be used, it lists many of the organizations and companies that support Frame Relay, and it includes when vendors claim Frame Relay will be available.

Frame Relay is the result of a number of converging networking market forces and technology trends. These forces include the need for greater transmission speed to support current applications; the increased use of LANs and the need to interconnect an enterprise's geographically separate LANs; the inherent burstiness of data traffic; the network world's multiplicity of protocols; and the need to minimize network costs.

At the same time, network technology has been moving in a similar direction. The widespread use of fiber optic-based transmission lines allows greater speed with fewer errors. This, combined with greater intelligence in network devices, implies that networks no longer need, for many applications, to provide many of the value-added services included in such feature-rich protocols or architectures as X.25. In addition, recent improvements have been made to ISDN data standards to improve its market acceptance.

If you combine the market drivers with the technology trends, you get the networking requirements that are fulfilled by Frame Relay.

Frame Relay is a very simple protocol. Like X.25, it is a DTE-DCE interface specification; it is not a network technology or architecture. It is based upon LAPD (Q.921), a layer 2 protocol. It uses only the "Core Aspects" of LAPD - framing, Frame Check Sequence checking, and addressing. There is no error recovery or frame retransmission - frames can be thrown away for any of a number of reasons. The framing and FCS are the same used for X.25 (HDLC), which means that current HDLC hardware can be used for Frame Relay. There is no defined layer 3 protocol. Most, if not all, implementations are based upon the "Consortium" specification.

At the physical layer, common interface examples are V.35 (at 56Kb to T1 speeds) to a co-located Frame Relay switch, and a 4-wire T1 to a Frame Relay carrier.

When compared to X.25, the Frame Relay DTE-DCE interface is much simpler, and has much less functionality.

Frames may be up to 8K octets in length, but the FCS is only useful in frames up to 4K octets in length. Addressing is specified by a 10-bit Data Link Connection Identifier (DLCI). Three bits in the header are used for congestion control

The DLCI is strictly local in significance to each Frame Relay interface on a network.

Addressing is different at each interface. This is similar to X.25 Permanent Virtual Circuits. Each DLCI specifies a connection between a pair of interfaces. 992 DLCIs are available at each interface to identify connections.

The Frame Relay specification includes a Local Management Interface (LMI), which is a simple protocol, loosely based upon ISDN, to perform signaling between the DTE and DCE. It is used by both the DTE and DCE to detect when the interface is up or down, and to provide information to the DTE as to which DLCIs currently exist and can be used to send data through the network.

The most controversial part of Frame Relay is its congestion management. The Consortium specification includes three bits. FECN (for forward explicit congestion notification to a destination interface), BECN (for backward explicit congestion notification to a source interface), and Discard Eligibility, which marks frames that may be thrown away first to avoid congestion. The ANSI specification is a superset of the Consortium specification; it also includes "implicit" congestion detection (from upper layer protocols, such as TCP) and a number of other mechanisms.

There are a number of optional features in the Consortium specification that not all vendors are implementing. These include global addressing, where every interface is addressed using the same DLCI from every other interface; multicasting, which sends frames to a number of other interfaces in a multicast group; and a number of LMI enhancements.

Frame Relay is implemented in the network using a number of different technologies, usually dependent upon the existing underlaying network architecture used by each vendor. All require fast, reliable transmission media, such as fiber optics, and some may require internal network overengineering to prevent congestion and packet loss.

Standards activities for Frame Relay are very active. It is originally based upon the IDSN I.122 and Q.921 standards. ANSI committee T1 has taken the lead, and has produced four standards in various stages of acceptance. The CCITT will probably accept the ANSI standards with very few changes. The biggest issue for the standards activities is resolving the differences between the Consortium and ANSI specifications. There are also a number of other outstanding technical issues that need to be resolved during the standards activity.

Frame Relay will typically be used for LAN interconnection and in wide-area networks to replace existing direct switch-switch trunks (especially T1 trunks). Frame Relay can lower network costs by replacing a number of direct trunks in a mesh network, or by allowing more optimal use of trunks between T1 multiplexers. It also removes the potential "milking machine" problem of many interconnections between collocated equipment, by multiplexing a number of network trunks over one physical interface. It also provides the advantage that every network DTE is logically a neighbor of every other DTE. There are a number of user concerns with Frame Relay that must be answered before its use will become widespread.

Frame Relay is being supported by the standards bodies and by a large number of vendors. The vendors have replaced the original "Consortium" with the Frame Relay Implementor's

Forum.

Public Frame Relay service is being quoted by services providers as being available as early as June 1991. No tariffs have yet been announced. Many vendors have announced that their DCE equipment is now available for purchase.

Frame Relay still has a number of questions that need to be resolved.

Who, What, When, Why, and How of Frame Relay

St. Louis IETF

Andrew G. Malis
BBN Communications

Topics

- Why is Frame Relay?
- What is Frame Relay?
- How will Frame Relay be used?
- Who is supporting Frame Relay?
- When will Frame Relay be available?

Why is Frame Relay?

- Frame Relay is the result of converging wide area networking market forces and technology trends
- WAN market forces:
  - The increased need for speed
  - Emergence of LAN-WAN internetworking and LAN interconnection
  - "Bursty" data communications
  - Multiplicity of protocols, and need for protocol transparency
  - Conservation of $$$

The Need for Speed

- Data network applications require more and more speed as they become more sophisticated
- Modern applications (esp. financial, insurance, CAD-CAM, and medical) are increasingly based on graphics and image transmission, rather than text - need to transmit many more bits
- End-to-end delay is an increasing problem - can be reduced by increasing bandwidth
- Distributed computing and bulk file transfer also require sending large amounts of data
- Traditional solution is using T1 lines, but this is expensive and inflexible

LAN Interconnection

- Most computing is moving from terminals and mainframes to PCs, workstations, and servers on LANs (client-server computing)
- Need to replace existing terminal-mainframe networking with LAN-LAN interconnection
- Need to interconnect an enterprise's LANs
- Users come to expect LAN-type speed, even over wide area nets
- Expensive to widely deploy high-bandwidth LAN-LAN links, but 64Kb doesn't make it anymore - too much transmission delay

Burstiness

- Data communications is inherently "bursty"; not a steady stream of data like video or audio
- Allows advantageous use of statistical multiplexing
- Burstiness and dedicated bandwidth (such as an entire T1 line) are a bad match - the bandwidth is required to reduce transmission time, not for capacity, so lines are largely underutilized (20% or less on average)
- In private T1 multiplexer networks, using fixed TDM slots also inefficient for data traffic

Multiplicity of Protocols

- Many, many networking protocols out there:

  TCP/IP, X.25, ISO, DECNET, SNA, SDLC, Novell IPX, XNS, proprietary router and PSN trunk protocols, LAN bridging, etc. etc.
- Need to be able to interconnect geographically separate networks transparently without protocol translation - expensive to develop and to use

Conservation of $$$

- These are lean times
- Transmission lines typically represent 80%-90% of total network costs
- Network managers need to get the most for their networking $$$, especially on these recurring costs

WAN Technology Trends

- WAN technology trends:
  - Widespread use of fiber optics transmission lines, allowing higher speeds and fewer transmission errors
  - Increased use of intelligent network end devices (PCs, workstations, and servers)

    => Reduced need for X.25 value-added services (error correction, guaranteed delivery)
  - Standardization and adoption of ISDN protocols, ISDN improvements for data to counter poor ISDN market acceptance

Combining the Two

- Combine the market drivers with the technology trends, and you get the following requirements for WAN data services:
  - High speed and low delay
  - Statistical multiplexing to support bursty traffic
  - Useful for LAN interconnection, supports lots of LAN protocols
  - "Lean & mean", especially when compared to X.25
  - Based upon and compatible with ISDN standards
  - Cheaper to use than dedicated circuits

What is Frame Relay?

- The simple description:
  - Start with X.25
  - At layer 2, keep same basic framing and FCS (so current hardware still works). Add addressing and a few control bits. Throw out retransmissions, link establishment, windows, error recovery. In other words, make it SIMPLE.
  - Completely throw out layer 3
  - Add in some simple interface management
  - The result is Frame Relay

With Slightly More Detail

- Like X.25, Frame Relay is a DTE-DCE interface specification
- It is NOT a network technology specification or architecture
- Based upon LAPD (Q.921), the ISDN version of LAPB (X.25 L2)
- It only uses the "Core Aspects" of LAPD - framing, FCS checking, and addressing bits
- There is no defined layer three protocol inside the frames, just data; any protocol whatsoever can be in the frame
- Common specification produced by "Consortium": StrataCom, cisco, Digital, Northern Telecom

Frame Relay Definitions

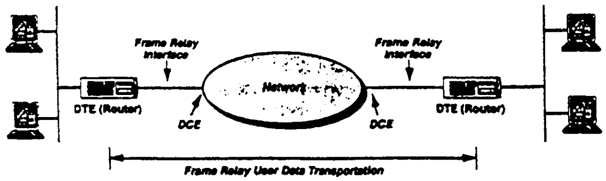

Physical Interface Examples

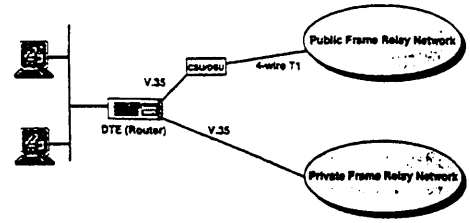

X.25 DTE-DCE Interface

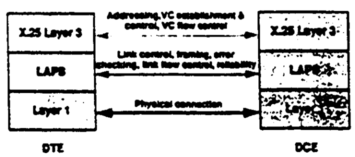

Frame Relay DTE-DCE Interface

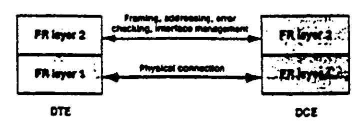

Data Link Interface

- Based upon "Core Aspects" of LAPD
- Provides framing, error checking, addressing bits
- Frames up to 8K octets in length; however, 16-bit FCS only useful for frames 4K or less in length
- Addressing: Data Link Connection Identifier (DLCI) - 10 bits
- Three bits used for congestion control:
  - Forward Explicit Congestion Notification (FECN)
  - Backward Explicit Congestion Notification (BECN)
  - Discard Eligibility

Frame Format

8	7	6	5	4	3	2	1	Octet
Flag 0	1	1	1	1	1	1	0	1
Data Link Connection Identifier (msb)						C/R	EA 0	2
DLCI (lsb)				FECN	BECN	DE	EA 1	3
User Data Field (Integral Number of Octets)								. . .
Frame Check Sequence (msb)								N-2
Frame Check Sequence (lsb)								N-1
Flag 0	1	1	1	1	1	1	0	N

Addressing

- DLCI contains addressing information
- Strictly local in significance
- Identifies a Permanent Virtual Circuit (PVC) between a pair of Frame Relay ports
- 10 bits in length
- 992 DLCIs available for PVCs
- One used for Local Management Interface
- One reserved for future SVC setup use
- Rest reserved for future use or options

Addressing Example

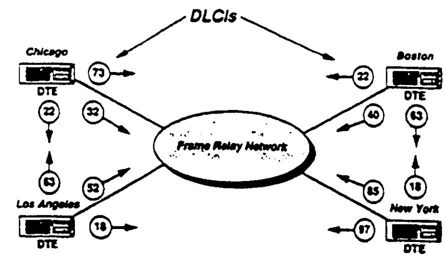

Local Management Interface (LMI)

- Loosely based on ISDN
- Uses DLCI 1023
- Notifies DTE of active PVCs in use, failed PVCs, new PVCs
- DTE polls DCE for current PVC status
- Heartbeat Polling and Keep Alive Sequencing are used to determine the state of the interface

Congestion Management

- Very controversial - basis of major criticisms of FR
- "Consortium" approach
  - FECN - forward explicit notification to destination
  - BECN - backwards explicit notification to source
  - DE - Discard Eligibility
- ANSI approach
  - Superset of "Consortium" approach
  - "Implicit" congestion detection (i.e. TCP)
  - Committed Information Rate, committed and excess burst sizes

Optional Features

- Global Addressing - each FR interface has a unique network-wide DLCI (limit of 992 interfaces/network)
- Multicasting - Set of FR interfaces is defined to be a "multicast group"
  - All receive frames sent to the group DLCI
  - Delivery is not guaranteed
- LMI enhancements
  - Asynchronous notification of new and failed PVCs
  - XON/XOFF flow control per PVC
  - Multicast DLCI status

Global Addressing

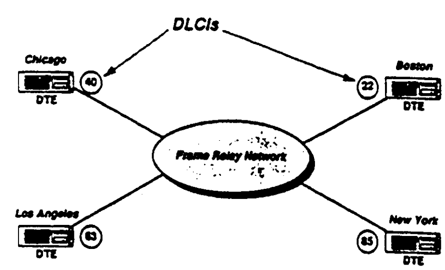

Frame Relay

How is Frame Relay Implemented?

- Given existing X.25 DCE interface, Frame Relay DCE interface is easy to add on (software change only)
- Internally, wide variety of technologies
  - T1 Mux (TDM) hybrid
  - Virtual circuits/frame switching (lite X.25)
  - ATM/Cell Relay
  - Datagram switching
  - All depend on fast, reliable transmission media (e.g. fiber optics)
  - May require overengineering to prevent congestion & packet loss

BBN Communications Page 25

Frame Relay

Frame Relay Standards Activity

- Originally based upon ISDN standards
  - I.122 - Framework for Advanced Packet Mode Services
  - Q.921 - ISDN Data Link Layer Specification
- US has taken lead with work in ANSI Committee T1
  - T1.606 - Frame Relay Service Description
  - T1.606 Addendum 1 - Congestion Management
  - T1.6ca - Protocol Definition
  - T1.6fr - Signalling Specification
- Vendors produced "Consortium" specification

BBN Communications Page 26

Frame Relay

Why Are Standards Important?

- International marketing
  - Agreement between ANSI and CCITT specifications
- Government focus on standards
- User focus on standards
- Benefits from standards
  - Public review and comment
  - Consumers want to interoperate
  - Vendors want market to expand
- Major issue: differences between "Consortium" and ANSI specifications

BBN Communications Page 27

Frame Relay

Status of Standards

Description	ANSI Standard Status		CCITT Standard Status	
Service Description	T1.606	Standard	I.2xy	Frozen
Congestion Management Strategy	T1.606 Adden.	Passed 1st Ballot	I.3xy	Frozen
Core Aspects	T1.6ca	Passed Let. Ballot	Q.922	Draft
Access Signalling	T1.6fr	Passed Let. Ballot	Q.93x	Draft

BBN Communications Page 28

Frame Relay

Other Technical Issues

- Control of changes to the "Consortium" specification
- Congestion management
- Extended addressing
- LMI specification, optional features
- Bringing specifications into agreement
- Standard for upper layer protocol encapsulation
- Certification & interoperability testing
- Network-network services and interconnection?
- Standardized network management interface?
- Access speeds above T1?
- SVCs?

BBN Communications Page 29

Frame Relay

How Will Frame Relay Be Used?

- Basic application - LAN interconnection over wide areas, especially for large networks
- Useful to provide PSN - PSN trunking, to replace direct trunks
- Useful for bulk data transfers, distributed computing

BBN Communications Page 30

Frame Relay

Before Frame Relay (1)

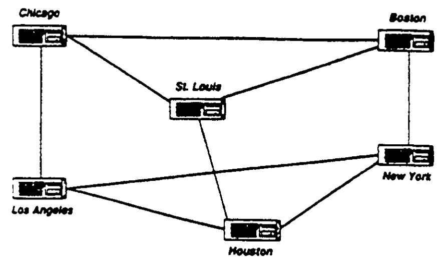

BBN Communications Page 31

Frame Relay

Before Frame Relay (2)

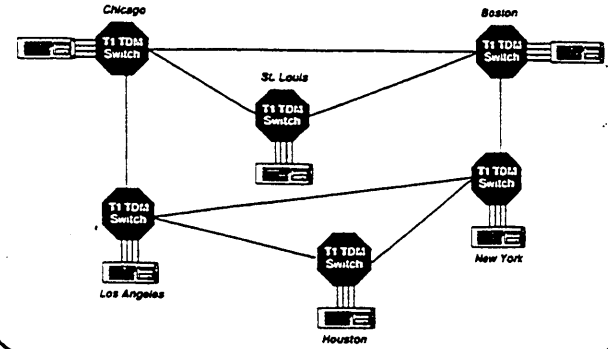

BBN Communications Page 32

Frame Relay

After Frame Relay

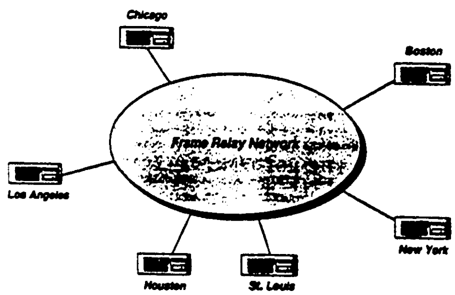

BBN Communications Page 33

Frame Relay

Logical Full Mesh

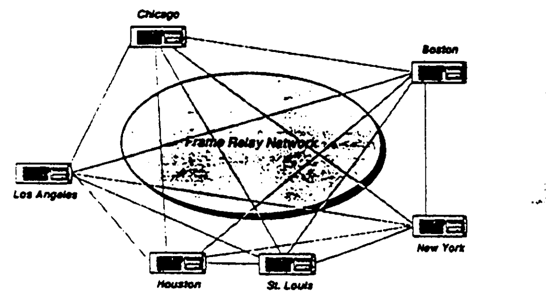

BBN Communications Page 34

Frame Relay

User Areas of Concern

- Perceived and actual value vs. private lines (tariffs)
- When will public service be available?
- How reliable will it be?
- Will there be enough bandwidth for each customer?
- How will congestion be managed?
- When will standards finalize?
- How much vendor support?
- Will BISDN/ATM/"fast packet" technologies overtake Frame Relay?
- What will market acceptance be like?

BBN Communications Page 35

Frame Relay

Who Is Supporting Frame Relay?

- Standards bodies:
  - ANSI
  - CCITT
  - IETF
- Frame Relay Implementor's Forum (was Consortium)
  - Industry business league
  - Technical input to standards bodies
  - Market development
  - Interoperability testing & certification

BBN Communications Page 36

FRIF Members

- 43 companies represented at first meeting
- Major companies:

3Com	Frame Relay Technology	Proteon
AT&T	Hughes	Racal-Milgo
BBN Communications	Infotron	StrataCom
Bell Northern Research	Netrix	Sync Research
BellSouth	Network Equipment Technologies	
cisco Systems	Network Systems	Tekelec
Codex	Newbridge	Telematics
CompuServe	Northern Telecom	Timeplex
Digital Equipment	NYNEX	US Sprint
Eastman Kodak	Pacific Bell	Vitalink

When Will Frame Relay be Available?

- Public service
  - US Sprint

    3Q91, 200+ points of presence (at least one in every LATA), Tokyo, London; 56Kb, 64Kb, & T1 access
  - BT Tymnet

    6/91, 56Kb or 64Kb access, T1 access in 1994
  - CompuServe

    3Q91, 56Kb access, future T1 access "based on user demand"
  - Geisco (no firm details)
- Tariffs have yet to be announced

When Will Frame Relay be Available?

- Private network vendors
  - StrataCom, NT, US Sprint, Hughes: "Now"
  - Codex, Newbridge, Netrix: 1Q91
  - Others also have products due "soon"

Conclusion

- Frame Relay may become the WAN data network technology of choice for the early 90s
- Logical bridge from dedicated lines and X.25 to ATM/SMDS/BISDN and other "fast packet" technologies
- Success primarily depends on competitive pricing and user acceptance
- Success secondarily depends on implementation details (speed, quality of service)
- Market differentiation areas for private network vendors are equipment pricing, optional features, and implementation details
- Still many open questions to resolve

Chapter 7

Distributed Papers

THE INFORMATION ON THE FOLLOWING PAGES

WAS ORIGINALLY PRESENTED BY NOEL CHIAPPA

AT THE 19TH INTERNET ENGINEERING TASK FORCE

DECEMBER 3 - 7, 1990

BOULDER, COLORADO

Internet Draft Noel Chiappa Consultant

The IP Addressing Issue

0 Status of this memo

This draft document will be submitted to the RFC editor as an informational Document. Distribution of this memo is unlimited. Please email comments to jnc@lcs.mit.edu or fax comments to (804) 898-7663.

1 Introduction

The packet layer of the IP architecture is about to enter a period of stress caused by deficiencies in the IP address. This stress is caused by a number of inter-related problems. This note describes these problems, lists some suggested solutions, and discusses pros and cons of each of those solutions.

2 The Problems

The IP address structure is facing three problems; some widely appreciated, and some not.

2.1 Exhaustion of Class B Address Space

The first, and one that has received a fair amount of notice in the technical community, is that the supply of class B network numbers is being used up quite rapidly. Around 16 thousand of these network numbers exist in the current address structure; xxxx thousand have already been assigned. Moreover, a simplistic statistical analysis indicates that the rate of growth closely fits an exponential curve, and if that curve continues to be followed, the existing stock of class B addresses will be used up in yy months.

Of course, the curve is really an S-curve, and will top out at some point, but it is not clear when. While the spread of the Internet in the U.S. academic community is slowing as coverage becomes ubiquitous (i.e. it is itself an S-curve), there are two largely untapped population factors which will keep the growth up for a while to come. The first is the spread of the Internet within the U.S. government and within U.S. companies; while there is currently some coverage in these groups, it has not yet reached the levels of the academic community, but can be expected to do so over time. The second is the spread of the Internet to the international community, where this three phase growth pattern will probably repeat.

2.2 Exhaustion of Entire IP Address Space

This is a short term version of a longer range problem, which is that the entire IP address space is being used up. At first glance, this seems unlikely, since the 32 bit space allows for approximately 4 billion distinct addresses.

However, a great deal of structure is imposed on this 32 bit field, in an attempt to allow agents inspecting an address to discover something about where in the network that address is. Imposing this structure does not allow the address space to be used efficiently, since there are many unused bit combinations. For example, if it is known that all addresses that start with the bit pattern 00010010 are at MIT (and all packets whose destination address starts with that pattern are sent to MIT), then unless there are 2^24 hosts at MIT (which there are not), all unused combinations starting with that pattern are effectively wasted. There is no way to avoid this without a fundamental change to IP addressing and routing.

It is unclear when the IP address space as a whole will be used up, assuming we continue in something like the current fashion. The IP address space is divided into different classes, and the usage rate in the various classes differs. As mentioned above, one class is already in some danger of running out. If the limit above is not to be the finale, presumably some reallocation of classes would be neccessary, but how long this will give us depends on how complete the reallocation is (more efficient schemes are more complex), as well as future growth patterns, which are not well understood.

Even if the address space were used efficiently (i.e. all possible bit combinations were valid addresses), if the Internet continues to grow and is not supplanted by an alternate networking technology, there will eventually be more hosts that can be addressed in this many bits (unlikely as this may seem at the moment).

2.2 Inutility of the Current IP Address Strucuture

Finally, and most important, the routing mechanisms themselves are breaking down due to the sheer number of different destinations they are required to track. This is in fact a far more serious problem (both in terms of immediate failures in the network as well as long range intractability) than the two above, although it may not be as well known or as apparent to the broader community.

The difficulty stems from the fact that the Internet is effectively a single level heirarchy. (For the purposes of this analysis, subnets can be ignored, although they did help a great deal.) The objects in terms of which routing is done are IP networks; each network is visible throughout the entire Internet (modulo local policy controls), and traffic cannot be routed to a destination in a network unless that network and something about its location are known to the source of the traffic and points in between. Simply examining a network number (as opposed to an entire address) will tell you nothing about where in the network that network is.

The problem is that any routing technology is limited in the number of discrete destinations it can track. The limits come in storage space to keep databases and processing time to do computations based on them, as well as bandwidth usage to send information around. Even systems which cost O(zzzzz) in the number of destinations (and most routing technologies aren't even that good) will have limits. The exact number of destinations at which a particular routing technology tops out differs, but as a general proposition the best systems of which we know now seem to have limits in the low tens of thousands with the current levels of hardware performance.

The solution to the problem is to adopt an address with more structure in it; this will allow groups of networks (and eventually groups of groups) to be represented as single objects for the purposes of routing, which will reduce the costs of routing to a manageable burden. Once again, there is no way to handle this without a fundamental change to IP addressing and routing; better hardware will not solve the problem as the size of the Internet is increasing faster than the capabilities of the hardware.

As a side note, the existing routing technology in the Internet has a very short useful lifetime. The problems of EGP are well known; in addition to not being able to handle cycles in the topology, since it sends the entire routing table as a single packet, it has a practical limit in the low thousands of networks. BGP will remove these two limits, but has little support for real policy routing, and in any case is subject to the ultimate size limits above inherent in the existing IP addresses. As far as IGP's go, the new link state protocols OSPF and dual IS-IS are very good protocols, but have the same problems as BGP. For example, OSPF can probably barely handle (if at all) the 2^14 class B network numbers currently allowed (the lack of definity is because the actual limits depend on the speed and memory size of the router), so increasing the number of class B networks by a factor of 10 (assuming this were somehow possible within the IP address space) would not really help.

3 Possible Solutions

A number of solutions have been proposed to solve the 'IP addressing problem'. However, most of them do not fully address all three of the problems outlined above. In some cases, this is deliberate; the proposals were put forward as temporary fixes to allow more time to be spent on a more complete solution; in other cases, the author was perhaps not aware of the full dimensions of the problems.

Two self-evident conclusions can be drawn from the list of problems above. First, any proposal which does not solve all the problems listed above is not a real solution, but simply a temporary patch, and is only worth considering if the extra time provided is needed and worth the cost.

Second, given the interrelation between addressing and routing, a more satisfactory solution almost certainly lies in considering and solving problems in the two areas together. Given that routing is by far the harder problem, the address structure chosen should be designed in light of the requirements of the final routing architecture; put another way, the address structure should be designed to make the job of the routing as easy as possible. Different address architectures can make a great deal of difference in the difficulty of routing them; to design an address structure without reference to the routing system that will provide the paths for the traffic is most unwise given the extreme technical challenges posed by the current requirements on routing in the Internet.

In any case, it is worth going down the list, explaining each, and listing the pros and cons of each. It is unfortunately not possible to go into detail on each solution, since that would require a long paper on each one. Note that most of these only address the address space limitations, not the routing problems, which, as have been pointed out, are actually the most severe.

In general, most solutions propose the adoption of some sort of new address; schemes vary as to how (or if) the new address is to be carried in the packet. Some propose to keep the existing IP packet format, while others propose to modify it.

Note also that most of the proposals are not mutually exclusive; one can take parts of one and mix it with parts of another to provide more complete solutions. For example, some of the ideas in the last proposed solution (such as the completion of Open Routing to handle new and more complex addresses and use of the existing 32 bit address as a UID) could obviously be used in other proposals, and several of the other solutions appear as pieces of or options in the last proposal.

3.1 Increase the Number of Class B Networks

The first (and only non-radical) proposed solution is to increase the number of class B network numbers available, either by allocating half the class A numbers to class B, or part of the class C address space.

The advantages of this are that it will require no changes in most hosts, and will extend the life of the current addresses substantially, since class B is the only one nearing exhaustion.

The disadvantages are that it is only a short term solution, it is not a solution to the most pressing problem (the third) anyway, it is not a solution to the address space exhaustion, it will require changes to all routers, and, if the class A subvariant is used, to some hosts (which persist in ignoring the requirements in the Host Requirements and understand the structure of IP addresses).

Nonetheless, this solution might be of use as an interim measure if more time is needed to implement a final plan.

3.2 Reformat the Existing IP Address

The second proposed solution is to redefine the interpretation of the existing 32 bit field to make it more useful. The primary proposal here is the 'bottoms up' proposal. Briefly, this proposal contemplates redefining the interpretation of the IP address to allow a multi-level heirarchy in which (through use of masks and assignment of extra bits to each level as needed) each level can grow efficiently, as needed. A way to perform policy routing using this address interpretation is also provided.

The advantages are that this does make efficient use of the address space, and does not require changes to hosts and many routers. It also improves the routing situation.

The disadvantage is that this particular proposal does not remove the ultimate limit on the size of the IP address; the proposal as written uses the existing 32 bit addresses. (Clearly, one could vary this to use new addresses, in which case it would fall into one of the classes below.)

3.3 Make IP Addresses Non-Globally Unique

The third proposed solution is to change IP so that addresses are not globally unique, but only unique in a single AS. This effectively corresponds to the scenario seen in a number of procotol families where distinct catenets with overlapping address spaces are glued together. This is usually undertaken to avoid renumbering an entire catenet when two catenets which developed separately are joined, rather than to expand a single catenet, but the details are the same whatever the motivation. In any case, experience in other protocol families with this solution is a useful guide.

The advantages are that it requires no change to the hosts, and it also avoids changes to any non-border routers.

The disadvantages, depending on the subvariant, are that either an overall size limit still remains, or a mechanism equivalently complicated to that of some of the more complete proposals must be developed to do the routing among AS's, etc. (This is an example of the adage that a problem swept under the carpet will always pop up somehere else.)

In one subvariant (seen where distinct catenets with overlapping address space assignments are joined, but probably not useful to the Internet), distant AS's are permanently mapped into the address space of the local AS, but in a place different from their 'native' address. Thus, both the third (routing difficulties) and second (total size limit) problems still exist.

In a different subvariant, either not all AS's can be mapped in, or external destinations are dynamically mapped into the address space of the local AS. In this case, a larger (and probably more complex) address must be adopted for use between AS's, routing must be designed to route it, and mechanisms developed to do the dynamic mapping. Depending on the details, this might look very similar to the next scheme.

In yet another subvariant, the AS's are joined at their edges not by packet level routers, but by application level gateways. This is in a sense a variant of the one above, except that the more complex address is the host name. The difficulty here is that the applications must in most cases be modified to pass the identification of the ultimate destination on start-up (in-band, since there is no out-of-band channel in TCP); very few (such as electronic mail) currently do this.

3.4 Define a New IP Address

The fourth proposed solution is to define a new IP address, which would be carried in an IP option in existing IP packets, perhaps with a pointer to the location of the option carried in a class E address. In one variant of this idea (similar to the one above) the AS number would be the extension (so that the existing IP address is again not unique across the entire Internet), and would be carried in an option. In another, the 'bottoms-up' addressing scheme would be used.

The advantages are that this removes the limits on the number of addresses, and this allows a solution to the routing problem, although one is not proposed, except in the 'bottoms-up' scheme.

The disadvantage is that all the hosts must be modified to generate the addresses in this fashion, and the overall size of the new address (assuming it is a multi-level address to make routing easier) is limited by the maximum amount of free space in the IP header. (If the hosts are not modified, but the new address is added by some agent, then this solution turns into the previous one.) An additional disadvantge in the case where the AS number is the extension is that this still only provides one extra level of heirarchy, which in the long run will not be enough.

3.5 Define a New IP Packet

This is a more radical attack on the problem, and since a 'clean slate' is available, the proposals differ substantially. The chief advantage of this approach is that other problems with IP can be solved at the same time, but these are outside the scope of this discussion. The chief disadvantage is that all the hosts must be modified to generate the new format packets, but some schemes include interoperable transition plans to ease this.

A number of proposals include this step as an option, to make the new addresses less of a kludge, or to provide extra capabilities.

3.6 Integrated New Host Identifier, Address and Routing.

The sixth proposal is a multi-stage solution which attacks a number of problems at once. It is in some respects related to the two above, since it contemplates a new IP address and (eventually) a new packet format, but it differs from them in introducing a new concept into the IP architecture (the Host Identifier) as well as tight coupling to a routing architecture.

It envisions creation of a new IP address, of varying length with a varying number of levels, upgrading the Open Routing protocol to handle these new addresses, conversion of interpretation of the existing 32 bit addresses to UID's for Host Identifiers, and, in the long run, a final step to allow the system to contain more than 2^32 distinct nodes. One possibility for the latter is to make the UID's locally significant only (using some of the mapping techniques laid out in the second sub-variant of 3.3); the other is a new IP packet format with 64 bit UID's.

The advantages of this proposal are varied. First, the structure of the new addresses can be oriented toward the main goal of making the routing easier. Second, a UID means that a number of problems with machines with more than one address or changing address can be attacked. Third, the initial retention of the existing address as the UID means that hosts and routers do not have to be changed right away, and the change to hosts is small; the address can be constructed by the first 'new-style' router (although eventually for efficiency the hosts should do this directly). Fourth, the existing 32 bit space can be used with maximal efficiency, delaying the date at which the exhaustion of this space must be tackled. Fifth, if the 64 bit UID path is chosen, and the phaseover started before the 32 bit space is used (so that the UID of any host in the new 64 bit space is simply the zero extension of its UID in the 32 bit space), there will not be any cases of new and old style hosts which cannot communicate due to the inabilty of the old address space to name hosts in the new space, which is the usual cause of problems in conversions.

The disadvantages of this proposal are that the first 'new-style' router will have to determine (and add, or otherwise retain) the address which corresponds to that destination UID, which is a repeated small cost, and in the long run all the hosts will have to be changed.

3.7 Use ISO

There are a number of sub-variants in this option. One possibility is to use just the ISO address; this is effectively a new address scheme, as described in 3.4, but the address would be in common with ISO. This would allow 'packet wrapping' in a simple algorithmic way (since no complex tables would be needed to map addresses), but it is not clear if this is useful. Packet translation would also be possible, with the same caveats.

Another is to use the ISO packet layer, but retaining the existing stream and above protocols; this is effectively a new packet format, as described in 3.5. This would unify the two catenets at the packet level, but this is probably not a big advantage given the multiprotocol router technology. Hosts from different suites would still not be able to interoperate without application or other gateways.

A third is to keep the TCP suite of applications, but to run them above the ISO stream protocols. This is little different from the previous scheme in its effects, but would have more far-reaching effects in terms of host software.

The last is a complete conversion to the ISO suite. This would have problems during the (likely lengthy) transition which are identical to the ones we see now with interoperation; service gateways are an imperfect (and in some ways crippling) solution, and general translation has proven impractical.

4 Conclusions

It is possible to draw some initial conclusions as to which of the possible classes of solutions is to be preferred.

To begin with the last, a conversion to ISO, while alluring, is not currently a useful option, for reasons both political and technical.

Reflection in both camps on the complex political situation between supporters of the two protocols has led to a strategy that actually appears to have some advantages. Basically, both sides agree that the overall goal is to create the best possible packet data architecture. Given that, a plausible case can be made that that end is actually better served by the existence of two competing efforts, provided that no energy is wasted in fruitless political combat between adherents of the two camps. As long as all the effort is directed towards improvement of the two protocols suites, the end result will be better than the result of a single effort, especially if each side feels free to inspect, learn and borrow from the work of the other.

In the context of this strategy, and also with reference to the technical status of the ISO family, adoption of an ISO conversion strategy to answer the problems of the Internet is clearly inappriate. Both protocol families face problems in handling growth, especially given the choice of the ISO designers in making administrative concerns the main spring of their address design, rather than topology and abstraction. The continuing existence of the two different standards (down to the address/packet layer, where the real problems lie), each with radically different means of creating designs, is necessary to create the best answer to these challenges. In addition, the ISO architecture is still lacking some necessary pieces (such as an inter-domain routing protocol) for use in replacement of the Internet.

As noted above, the first two options are not really suitable. Since neither removes the straightjacket of the existing IP address (in length, if not in structure), any solution based on them would be short term at best.

Ruling out the third option (and all its subvariants) is a little harder to rigorously justify (since strictly speaking they can meet all the requirements), but it seems clear in the light of engineering experience that the existence of a single global system in which each host has a unique tag is extremely valuable and robust. Previous experiences with mapping solutions indicate that while they can eventually be made to work, they are inferior in many ways. Many other goals (such as security) are much easier to tackle in such a system. The last subvariant (with application level gateways) will crimp development and deployment of new applications, since lack of direct packet level connectivity will require the creation of application gateways as well as the applications themselves before a new application can spread. One of the chief advantages of the IP system (as compared with the various local solutions arrayed around the NCP-based ARPANet which preceeded it) is the direct packet access to all corners of the network.

The choice thus comes down to the middle three options of a new address, packet format, or the integrated rework of the packet layer. This is not really a choice between differing approaches, but simply a choice of how expansive a rework is desired or feasible.

Given that major changes are going to be necessary in any case, and given that a conversion/interoperation plan with minimal up-front costs is available, it seems likely that the best course for the long-term is the third; a complete and integrated rethinking and reworking of the basic addressing and routing facet of the packet level of the architecture.

IP ADDRESS SPACE PROBLEMS

1 QUESTION

3 PROBLEMS

7 SOLUTIONS

1 QUESTION

WHAT IS TCP/IP'S LIFETIME GOING TO BE?

ASSUMPTION:

- *PACKET SWITCHING IS THE FUTURE OF DATA COMMUNICATIONS.*

POSSIBILITIES:

1 - *IP IS REPLACED BY ISO/SOMETHING ELSE QUICKLY.*

2 - *IP IS REPLACED BY ISO/SOMETHING ELSE SLOWLY.*

3 - *IP AND ISO CONTINUE TOGETHER INDEFINITLEY.*

4 - *IP TAKES OVER.*

CASE 1 IS UNLIKELY; THE INTERNET IS TOO BIG, AND ISO IS MISSING TOO MANY PIECES.

NO MATTER WHICH OF THE OTHER THREE HAPPEN, IP IS GOING TO BE AROUND FOR A LONG TIME.

ADDITIONAL VIEW:

THE COMMUNITY IS BEST SERVED IF POLITICAL CONFLICT BETWEEN COMPETING PROTOCOL FAMILIES IS AVOIDED. RATHER, EACH PROTOCOL SHOULD SERVE AS A MODEL AND INSPIRATION TO THE OTHERS, TO PROVIDE THE USERS IN THE END WITH THE BEST POSSIBLE SYSTEM.

CONCLUSION:

IP NEEDS TO DEVELOP THE BEST POSSIBLE SOLUTION IT CAN, WHETHER AS AN OPERATIONAL SYSTEM OR AS A MODEL FOR OTHER PROTOCOLS; THIS IS WHAT IS BEST FOR EVERYONE.

3 PROBLEMS

\# 1 CLASS B NETWORK NUMBERS ARE RUNNING OUT.

- *16,000 EXIST, 2,500 ARE ALLOCATED.*
- *AT THE PRESENT RATE, WE WILL RUN OUT IN MID 1994.*

\# 2 THE IP ADDRESS SPACE IS RUNNING OUT.

- *STRUCTURE WHICH IS IMPOSED ON THE 32 BIT SPACE FOR ROUTING PURPOSES MEANS IT IS USED VERY INEFFICIENTLY.*
- *THE EXHAUSTION OF EXISTING CLASS B SPACE MEANS WE WILL HAVE USED 25% OF THE ADDRESS SPACE THERE ALONE.*

\# 3 THERE ARE TOO MANY NETWORK NUMBERS.

- *THE SHEER NUMBER OF DIFFERENT DESTINATIONS IS OVERWHELMING THE ROUTING.*
- *THIS IS THE LEAST WELL-KNOWN PROBLEM, BUT IS IN FACT THE WORST, BOTH IN TERMS OF THE SHORT-TERM PROBLEMS AS WELL AS LONG-TERM INTRACTABILITY.*
- *THE SINGLE LEVEL HEIRARCHY OF NETWORK NUMBERS IS UNWORKABLE IN THE LONG RUN.*
- *THE LIMITS OF THE EXISTING SYSTEM ARE FUNDAMENTAL TO ALL ROUTING TECHNOLOGY, AND ARE IN LINE BANDWIDTH, STORAGE SPACE AND PROCESSING TIME. THE SYSTEM IS GROWING FASTER THAN THE HARDWARE IS BECOMING MORE POWERFUL.*

7 SOLUTIONS

\# 1 - CREATE MORE NETWORK NUMBERS.

\# 2 - USE THE EXISTING ADDRESS SPACE MORE EFFICIENTLY.

\# 3 - MAKE ADDRESSES NON-GLOBALLY UNIQUE.

\# 4 - DEVISE A NEW ADDRESS.

\# 5 - DEVISE A NEW PACKET FORMAT.

\# 6 - ADD AN ADDITIONAL ADDRESS.

\# 7 - USE ISO.

SOLUTIONS

NOTES:

- SOME SOLUTIONS DO NOT SOLVE ALL THE PROBLEMS; THEY WERE PROPOSED ONLY AS PATCHES.
- THE FINAL DESIGN MAY INCLUDE ELEMENTS OF MORE THAN ONE SCHEME; THEY ARE NOT NECESSARILY EXCLUSIONARY.

POINTS:

- ANY PROPOSAL WHICH DOES NOT ATTACK ALL THREE IS JUST A PATCH, AND ONLY WORTH CONSIDERING IF THE EXTRA TIME IS NEEDED AND WORTH THE COST OF THE PATCH.
- TO FIX THE THIRD PROBLEM, WE NEED AN ADDRESS WITH MORE STRUCTURE; THIS MEANS A FUNDAMENTAL CHANGE TO IP ADDRESS AND ROUTING.
- GIVEN THE INTERRELATED NATURE OF ADDRESSING AND ROUTING, AND THE EXTREME DIFFICULTY OF THE ROUTING PROBLEM, THE BEST SOLUTION IS TO MAKE THE ADDRESSES BE WHATEVER MAKES THE ROUTING EASIEST.

SOLUTION ONE

CREATE MORE NETWORK NUMBERS.

o INCREASE THE NUMBER OF CLASS B NETWORK NUMBERS.

o TAKE THE SPACE FROM EITHER THE END OF CLASS A, OR PART OF CLASS C.

+ REQUIRES NO CHANGE TO MOST HOSTS.

+ EXTENDS LIFE OF CURRENT ADDRESS STRUCTURE SUBSTANTIALLY, SINCE ONLY CLASS B IS CLOSE TO RUNNING OUT.

- REQUIRES CHANGES TO ALL ROUTERS.

- DOES NOT SOLVE PROBLEMS 2 OR 3.

o MAY BE USEFUL AS A SHORT-TERM PATCH.

SOLUTION TWO

USE THE EXISTING ADDRESS SPACE MORE EFFICIENTLY.

o USE THE ALTERNATIVE ALLOCATION OF THE EXISTING 32 BIT FIELD.

o "BOTTOMS-UP" PROPOSAL INCLUDES A MULTI-LEVEL HEIRARCHY WITH VARIABLE NUMBER OF VARIABLE LENGTH LEVELS.

+ REQUIRES NOT CHANGE TO MOST HOSTS.

+ MAKES EFFICIENT USE OF THE ADDRESS SPACE.

+ DOES SOLVE PROBLEMS OF ROUTING OVERLOAD.

- REQUIRES CHANGES TO ALL ROUTERS.

- DOES NOT REMOVE ULTIMATE LIMIT ON THE ADDRESS SPACE.

SOLUTION THREE

MAKE ADDRESSES NON-GLOBALLY UNIQUE.

+ REQUIRES NO CHANGE TO ANY HOSTS.

+ REQUIRES NO CHANGE TO MANY ROUTERS.

- AN EQUIVALENTLY COMPLICATED MECHANISM TO THE MORE COMPLEX PROPOSALS MUST BE DEVELOPED TO HANDLE TRAFFIC FLOWING BETWEEN AREAS.

o ONE VARIANT TEMPORARILY MAPS EXTERNAL DESTINATIONS INTO THE LOCAL ADDRESS SPACE.

- EXPERIENCE WITH MAPPING SOLUTIONS SHOWS THEY CAN BE MADE TO WORK, BUT ARE COMPLICATED AND LESS ROBUST.

o ANOTHER VARIANT USES SERVICE LEVEL GATEWAYS AT THE BORDERS.

+ DOES PREVENT SECURITY ATTACKS ON FAULTY APPLICATIONS.

- DEPENDING ON THE IMPLEMENTATION, APPLICATIONS MAY HAVE TO BE MODIFIED TO PASS THE ULTIMATE DESTINATION ON START-UP.

- NEW APPLICATIONS CANNOT BE WIDELY DEPLOYED UNTIL THE SERVICE GATEWAYS ARE DEPLOYED.

SOLUTION FOUR

DEVISE A NEW ADDRESS.

- o CARRIED AS AN IP OPTION IN EXISTING PACKETS.
- o A POINTER TO THE ADDRESS COULD BE IN A CLASS E ADDRESS.
- o ONE OPTION IS TO MAKE THE AS NUMBER AN EXTENSION TO THE ADDRESS.
- o THE "BOTTOMS-UP" STYLE OR SOME OTHER NEW ADDRESS MIGHT BE USED.
- + REMOVES LIMIT ON ADDRESS SPACE.
- + ALLOWS SOLUTION TO THE ROUTING PROBLEM.
- - REQUIRES CHANGES IN ALL HOSTS.
- - SIZE OF THE NEW ADDRESS WILL BE LIMITED BY THE IP HEADER SIZE.
- - ADDING AS NUMBERS ONLY PROVIDE ONE EXTRA LEVEL OF HEIRARCHY, NOT ENOUGH IN THE LONG RUN.

SOLUTION FIVE

DEVISE A NEW PACKET FORMAT.

- o SOME SCHEMES INCLUDE INTEROPERABLE TRANSITION PLANS.
- + ALLOWS SOLUTION OF OTHER IP PROBLEMS.
- + REMOVES LIMITS ON ADDRESS SIZE.
- - ALL HOSTS MUST BE MODIFIED (EVENTUALLY).

SOLUTION 6

ADD AN ADDITIONAL ADDRESS.

- o THE EXISTING ADDRESS WOULD BE RETAINED AND USED AS A HOST IDENTIFIER.
- o HAS A TRANSITION SCHEME TO 64 BIT HOST IDENTIFIERS.
- o WOULD USE A NEW ADDRESS, AND PROBABLY A NEW PACKET FORMAT.
- o TEMPORARY CONTINUED USE OF THE EXISTING ADDRESS ALLOWS AN INCREMENTAL DEPLOYMENT IN ROUTERS.
- + DOES NOT REQUIRE ANY CHANGE TO EXISTING HOSTS.
- + ALLOWS SOLUTION TO ALL THREE PROBLEMS.
- + HOST IDENTIFIER ALLOWS SOLUTION OF OTHER PROBLEMS, SUCH AS MOBILE AND MULTI-HOMED HOSTS.
- + THE EXISTING 32 BIT ADDRESS SPACE CAN BE USED WITH MAXIMAL EFFICIENCY.
- - ALL ROUTERS WOULD HAVE TO BE CHANGED.
- - FIRST HOP ROUTER WILL HAVE TO ADD NEW ADDRESS (OR CONVERT THE PACKET IF A NEW PACKET FORMAT IS USED).

SOLUTION SEVEN

USE ISO

- o ONE OPTION IS TO USE JUST THE ISO ADDRESS.
- + ALLOWS PACKET MAPPING.
- - HAS ALL THE DISADVANTAGES OF A NEW ADDRESS.
- - THE ISO ADDRESS IS NOT ORGANIZED TOPOLOGICALLY, SO IT WILL NOT BE A GREAT HELP IN SOLVING THE THIRD PROBLEM.
- o ANOTHER IS TO USE THE ISO PACKET LEVEL.
- + WOULD UNIFY THE TWO CATENETS AT THE PACKET LEVEL.
- - HOSTS FROM DIFFERENT SUITES WOULD STILL NOT BE ABLE TO INTEROPERATE.
- o KEEP THE IP/TCP APPLICATIONS, ON TOP OF TP4.
- - LITTLE DIFFERENT IN EFFECT FROM THE PREVIOUS, BUT DIFFERENT SOFTWARE RAMIFICATIONS.
- o CONVERT TO ISO COMPLETELY.
- - PROBLEMS DURING CONVERSION, SINCE SERVICE GATEWAYS ARE MISSING.
- - ISO SUITE HAS SCALING PROBLEMS, ETC., AS WELL.

CONCLUSIONS

- A CONVERSION TO ISO, WHILE ALLURING, IS IMPRACTICAL FOR TECHNICAL AND POLITICAL REASONS.
- CREATING MORE NETWORK NUMBERS, OR USING THE EXISTING SPACE MORE EFFICIENTLY, WILL NOT PREVENT A PROBLEM IN THE LONG RUN ANYWAY.
- SOLUTIONS TO NON-GLOBALLY UNIQUE ADDRESSES HAVE PROVEN TO BE LESS DESIREABLE IN PRACTICE.
- CHOICE COMES DOWN TO A NEW ADDRESS, A NEW PACKET FORMAT, OR ADDING THE ADDITION ADDRESS. THIS IS NOT REALLY A CHOICE BETWEEN DIFFERING APPROACHES, BUT SIMPLY A CHOICE OF HOW COMPLETE A REWORK IS DESIRED OR FEASIBLE.

o GIVEN THAT MAJOR CHANGES ARE GOING TO BE NECESSARY ANYWAY, THE BEST COURSE FOR THE LONG-TERM IS A COMPLETE REWORK.

Appendix A

Attendees

David Abbajay
cisco Systems
Network Management
1525 O'Brian Drive
Menlo Park, CA 94025
(W)415-688-8245
(WF)415-688-8282
dabbajay@cisco.com

Shukri Abdallah
University of Maryland
5401 Wilkens Avenue
Baltimore, MD 21228
(W)301-455-3000
(WF)301-455-3969
abdallah@umbc3.umbc.edu

Vikas Aggarwal
Princeton University
JVNCnet
B6von Neumann
Princeton, NJ 08544
(W)609-258-2403
(WF)609-258-2424
vikas@JVNC.net

Philip Almquist
Consultant
214 Cole Street
San Francisco, CA 94117-1916
(W)415-752-2427
(WF)415-752-2229
almquist@jessica.stanford.edu

William Anderson
The MITRE Corporation
Burlington Road
M/S EO66
Bedford, MA 01730
(W)617-271-8838
(WF)617-271-8953
wda@mitre.org

Kathy Atnip
Washington University
Office of the Network Coordinator
One Brookings Drive
Campus B, 1048
St. Louis, MO 63130-4899
(W)314-726-7387
(WF)314-726-7302
kathy@wugate.wustl.edu

Karl Auerbach
Sun Microsystems, Inc.
2550 Garcia Avenue
M/S MTV 14-49
Mountain View, CA 94043-1100
(W)415-336-2090
(WF)415-336-3156
karl@eng.sun.com

William Babson
Penril Datacomm Networks
1300 Quince Orchard Boulevard
Gaithersburg, MD 20878
(W)301-921-8600
(WF)301-921-8376
bill%penril@uunet.UU.NET

Douglas Bagnall
Hewlett-Packard
Apollo Systems Division
250 Apollo Drive
Chelmsford, MA 01824
(W)508-256-6600
bagnall_d@apollo.hp.com

Fred Baker
Advanced Computer Communications
Engineering
720 Santa Barbara Street
Santa Barbara, CA 93101
(W)805-963-9431
(WF)805-962-8499
fbaker@emerald.acc.com

Atul Bansal
Digital Equipment Corporation
550 King Street
M/S LKG1-3/L06
Littleton, MA 01460
(W)508-486-5355
(WF)508-486-7417
bansal@netrix.nac.dec.com

Ballard Bare
Hewlett-Packard
8000 Foothills Blvd.
Roseville, CA 95678
(W)916-785-5608
(WF)916-786-9185
bare@hprnd.rose.hp.com

William Barns
The MITRE Corporation
7525 Colshire Drive
M/S W425
McLean, VA 22102-3481
(W)703-883-6832
(WF)703-883-7142
barns@gateway.mitre.org

Richard Basch
Massachusetts Institute of Technology
Project Athena
1 Amherst Street
M/S E40-318
Cambridge, MA 02139
(W)617-253-0100
(WF)617-258-8736
probe@mit.edu

Charles Bazaar
Emulex Corporation
3545 Harbor Boulevard
Costa Mesa, CA 92626
(W)714-668-5416
bazaar@emulex.com

Bob Beach
Ultra Network Technologies
101 Daggett Drive
San Jose, CA 95134
(W)408-972-0100
bob@ultra.com

Jordan Becker
Advanced Network & Services, Inc.
Network Services
100 Clearbrook Road
Elmsford, NY 10523
(W)914-789-5315
(WF)914-789-5310
becker@nis.ans.net

Warren Benson
University of Nebraska at Omaha
6001 Dodge Street
Omaha, NE 68182-0051
(W)402-554-3470
(WF)402-554-3475
wbenson@zeus.unomaha.edu

David Benton
National Institute of Health
Nat'l. Center for Human Genome Rsch
9000 Rockville Pike
Building 38A, Room 610
Bethesda, MD 20892
(W)301-496-7531
(WF)301-480-2770
benton@bio.nlm.nih.gov

Arthur Berggreen
Advanced Communications Corporation
720 Santa Barbara Street
Santa Barbara, CA 93101
(W)805-963-9431
(WF)805-962-8499
art@acc.com

Ron Bhanukitsiri
Digital Equipment Corporation
305 Lytton Avenue
Palo Alto, CA 94301
(W)415-688-1528
`rbhank@wsl.dec.com`

Joe Blackmon
National Center for
Supercomputing Applications
605 E. Springfield Avenue
Champaign, IL 61820
(W)217-244-4019
(WF)217-244-1987
`blackmon@ncsa.uiuc.edu`

James Blaine
Electronic Radiology
Laboratory
Mallinckrodt Institute of Radiology
Washington Univ. Medical School
4524 Scott Avenue, Room 1150
M/S Box 8131
St. Louis, MO 63110

Nathaniel Borenstein
Bell Communications Research
Room MRE 2A-274
445 South Street
Morristown, NJ 07960
(W)201-829-4270
(WF)201-829-7019
`nsb@thumper.bellcore.com`

David Borman
Cray Research, Inc.
655F Lone Oak Drive
Eagan, MN 55121
(W)612-683-5571
(WF)612-683-5599
(W)612-452-6650
`dab@cray.com`

Steve Bostock
Novell, Inc.
Software Engineering
2180 Fortune Drive
San Jose, CA 95131
(W)408-473-8203
(WF)408-435-1706
`steveb@novell.com`

William Bostwick
Department of Energy
Federal Networking Council
M/S ER-7
Washington, DC 20545
(W)703-276-3533
(W)301-353-5800
(W)703-276-3526
`bostwick@DARPA.MIL`

Andreas Bovopoulos
Washington University
Department of Computer Science
Campus Box 1045
St. Louis, MO 63130
(W)314-726-4465
(WF)314-726-7302
`andreas@patti.wustl.edu`

Richard Bowles
Space Telescope Science Institute
3700 San Martin Drive
Baltimore, MD 21218
(W)301-338-4852
(WF)301-338-4767
`bowles@stsci.edu`

Helen Bowns
Bolt, Beranek, and Newman
Communications
87 Fawcett Street
Cambridge, MA 02138
(W)617-873-2411
(WF)617-873-4086
`hbowns@bbn.com`

David Bridgham
Epilogue Technology
2610 Monstrat Avenue
Belmont, CA 94002
(W)415-594-9268
(WF)415-594-1141
dab@asylum.sf.ca.us

Scott Brim
Cornell University
Cornell Information Technologies
143 Caldwell Hall
Ithaca, NY 14853-3801
(W)607-255-5510
(WF)607-255-5771
swb@devvax.tn.cornell.edu

Ronald Broersma
Naval Ocean Systems Center
Code 914
San Diego, CA 92152-5000
(W)619-553-2293
(WF)619-553-6811
ron@nosc.mil

Caralyn Brown
Wellfleet Communications, Inc.
15 Crosby Drive
Bedford, MA 01730
(W)617-275-2400
(WF)617-275-5001
cbrown@wellfleet.com

Howard Brown
Cabletron Systems, Inc.
P.O. Box 5005
35 Industrial Way
Rochester, NH 03867
(W)603-332-9400
(WF)603-332-5917
brown@ctron.com

Theodore Brunner
Bell Communications Research
445 South Street
M/S MRE 2P252
Morristown, NJ 07962-1910
(W)201-829-4678
(WF)201-267-9298
tob@thumper.bellcore.com

Christopher Bucci
Datability, Inc.
322 Eighth Avenue
New York, NY 10001
(W)212-807-7800
bucci@pluto.dss.com

Milind Buddhikot
Washington University
Computer Science Department
410 Bryan Hall
1 Brookings Drive
St. Louis, MO 63130
(W)314-726-4163
milind@dworkin.wustl.edu

Jeffrey Burgan
Sterling Software Federal Systems
Ames Research Center
M/S 233-8
Moffett Field, CA 94035
(W)415-604-5705
(W)415-964-9900
(WF)415-604-0063
jeff@nsipo.nasa.gov

Randy Butler
National Center for
Supercomputing Applications
605 East Springfield Avenue
Champaign, IL 61820
(W)217-244-8285
(WF)217-244-1987
rbutler@ncsa.uiuc.edu

Ross Callon
Digital Equipment Corporation
550 King Street
M/S LKG1-2/A19
Littleton, MA 01460-1289
(W)508-486-5009
(WF)508-486-5279
callon@bigfut.enet.dec.com

Lida Carrier
Apple Computer, Inc.
A/UX Engineering
10300 Bubb Road
M/S 50-UX
Cupertino, CA 95014
(W)408-974-8325
lida@apple.com

Eric Carroll
University of Toronto
Computing Services
4 Bancroft Avenue, Room 102
Toronto, Ontario M5S 1C1
Canada
(W)416-978-6134
(WF)416-978-6620
(W)416-961-7187
eric@utcs.utoronto.ca

Charles Carvalho
Advanced Computer Communications
720 Santa Barbara Street
Santa Barbara, CA 93101
(W)805-963-9431
(WF)805-962-8499
charles@sage.acc.com

Jeffrey Case
The University of Tennessee
Department of Computer Science
3001 Kimberlin Heights Road
107 Ayres Hall
Knoxville, TN 37996
(W)615-974-5067
(WF)615-573-1434
case@cs.utk.edu

Stephen Casner
USC/Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292-6695
(W)213-822-1511
(W)213-823-6714
casner@isi.edu

Peter Caswell
Pacer Software, Inc.
1900 West Park Drive, #280
Westborough, MA 01581
(W)508-898-3300
(WF)508-336-1356
uunet!pacvax!pfc

Vinton Cerf
Corporation for National
Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
(W)703-620-8990
(WF)703-620-0913
vcerf@nri.reston.va.us

Martina Chan
Motorola, Inc.
Information Technology Management
1299 E Algonquin Road
M/S IL01-03
Schaumburg, IL 60196
(W)708-576-1957
(WF)708-576-6388
mchan@mot.com

Cho Chang
Hewlett-Packard
Apollo Systems Division
250 Apollo Drive
M/S CHD-02-CN
Chelmsford, MA 01824
(W)508-256-6600
(WF)508-229-2150
chang_c@apollo.hp.com

Chris Chiotasso
Spartacus, Inc.
Micro Development Group
One Lowell Research Center
847 Rogers Street
Lowell, MA 01852
(W)508-937-1600
(WF)508-937-0445
chris@roswell.spartacus.com

Cyrus Chow
NASA
Ames Research Center
M/S 233-18
Moffett Field, CA 94035-1000
(W)415-604-6843
(WF)415-604-4003
cchow@orion.arc.nasa.gov

Paul Ciarfella
Digital Equipment Corporation
Communications Systems Engineering
153 Taylor Street
M/S TAY2-2/B4
Littleton, MA 01460
(W)508-952-3548
(WF)508-952-3023
ciarfella@levers.enet.dec.com

George Clapp
Ameritech Services
Science and Technology
Meadows Corporate Center
2820 West Golf Road
Rolling Meadows, IL 60008
(W)708-806-8318
(WF)708-806-8292
meritec!clapp@uunet.uu.net

Alan Clegg
Microelectronics Center
of North Carolina
Center for Communications
3021 Cornwallis Road
RTP, NC 27709-2889
(W)919-248-1120
(WF)919-248-1405
abc@concert.net

Graham Cobb
Digital Equipment Co. Ltd.
Worton Grange
Reading, Berks. RG2 0TU
UK
(W)+44 734 853917
cobb@marvin.enet.dec.com

Richard Colella
National Institute of
Standards and Technology
Room B217
Building 225
Gaithersburg, MD 20899
(W)301-975-3627
(WF)301-590-0932
colella@osi3.ncsl.nist.gov

Robert Collet
US Sprint
Government Systems Division
12490 Sunrise Valley Drive
M/S VARESB221A
Reston, VA 22096
(W)703-689-5133
(WF)703-689-5177
/pnrobert.d.collet/o=us.sprint/admd=telemail/c=us/@sprint.com=

Jill Colquhoun
Washington University
Office of the Network Coordinator
One Brookings Drive
Campus Box 1048
St. Louis, MO 63130-4899
(W)314-726-7388
(WF)314-726-7302
jill@wugate.wustl.edu

Rob Coltun
510 New Mark Esplanade
Rockville, MD 20850
(W)301-240-2168
(W)301-921-8600
rcoltun@trantor.umd.edu

Don Coolidge
Silicon Graphics, Inc.
2011 N. Shoreline Boulevard
Mountain View, CA 94039
(W)415-962-3374
(WF)415-969-2314
coolidge@speaker.wpd.sgi.com

Jerome Cox
Washington University
Department of Computer Science
415 Wicksworth Lane
St. Louis, MO 63122
(W)314-889-6132
jrc@wucs1.wustl.edu

Tracy Cox
Bell Communications Research
NVC 1C354
331 Newman Springs Rd.
Redbank, NJ 07701
(W)908-758-2107
(WF)908-747-3424
tacox@sabre.bellcore.com

Bruce Crabill
University of Maryland
Computer Science Center
Bldg. 224 Room 1313
College Park, MD 20742
(W)301-405-3056
(WF)301-314-9220
bruce@umdd.umd.edu

Chuck Cranor
Washington University
Computer Science Department
406 Bryan Hall
1 Brookings Drive
St. Louis, MO 63130
(W)314-726-4203
chuck@maria.wustl.edu

Mark Crispin
University of Washington
3737 Brooklyn Avenue NE
Seattle, WA 98105
(W)206-543-5762
(WF)206-543-3909
mrc@cac.washington.edu

David Crocker
General Magic, Inc.
444 Castro Street
Suite 1200
Mountain View, CA 94041
(W)(415) 965-0400
(WF)(415) 965-9424
david_crocker@palo-alto.pa.dec.com

Stephen Crocker
Trusted Information Systems
3060 Washington Road
Route 97
Glenwood, MD 21738
(W)301-854-6889
(WF)301-854-5363
crocker@tis.com

Nabil Damouny
NEC America
Transmission Development Division
110 Rio Robles
San Jose, CA 95134
(W)408-433-2956
(WF)408-433-2950

Megan Davies
Corporation for National
Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
(W)703-620-8990
mdavies@nri.reston.va.us

James Davin
Massachusetts Institute of
Technology
Laboratory for Computer Science
545 Technology Square
M/S NE43-507
Cambridge, MA 02139
(W)617-253-6020
(WF)617-258-8682
jrd@ptt.lcs.mit.edu

Steve Deering
Xerox Corporation
Palo Alto Research Center
Computer Science Laboratory
3333 Coyote Hill Road
Palo Alto, CA 94304
(W)415-494-4839
deering@xerox.com

Barbara Denny
SRI International
333 Ravenswood Avenue
M/S EK372
Menlo Park, CA 94025
(W)415-859-2812
(WF)415-859-4812
denny@sri.com

Zubin Dittia
Washington University
Bryan Hall, Room 412
CCRC
St. Louis, MO 63130
(W)314-726-4163
zubin@dworkin.wustl.edu

Kurt Dobbins
Cabletron Systems
Software Engineering
35 Industrial Way
Rochester, NH 03867
(W)603-332-9400
(WF)603-332-1019
dobbins@ctron.com

Sean Donelan
Data Research Associates, Inc.
1276 North Warson Road
St. Louis, MO 63132-1806
(W)314-432-1100
(WF)314-993-8927
sean@dra.com

Avri Doria
Clearpoint Research
Corporation
Research and Development
35 Parkwood Drive
Hopkinton, MA 01748
(W)508-435-0900
(WF)508-435-7504
avri@asherah.clearpoint.com

Ralph Droms
Bucknell University
Computer Science Department
323 Dana Engineering
Lewisburg, PA 17837
(W)717-524-1145
(WF)717-524-3760
droms@bucknell.edu

Martin Dubetz
Washington University
Office of the Network Coordinator
Campus Box 1048
One Brookings Drive
St. Louis, MO 63130-4899
(W)314-726-7389
(WF)314-726-7302
dubetz@wugate.wustl.edu

John Dudeck
California Polytechnic
State University
611 Couper, #6
San Luis OBispo, CA 93405
(W)805-545-9549
(WF)805-543-2136
jdudeck@polyslo.calpoly.edu

Bill Durham
McDonnell Douglas
P.O. Box 516
M/S 080 1480
St. Louis, MO 63166
(W)314-895-5296
(WF)314-895-5297
durham@MDC.COM

Tom Easterday
Ohio State University
Instr. and Res. Computer Center
Academic Computing Services
1971 Neil Avenue, Room 406
Columbus, OH 43210
(W)614-292-4027
(WF)614-292-7081
tom@cic.net

Jeffrey Edelheit
The MITRE Corporation
WCI Security Technical Center
7525 Colshire Drive
McLean, VA 22102-3481
(W)703-883-7586
(WF)703-883-1397
edelheit@smiley.mitre.org

Nadya El-Afandi
Network Systems Corporation
Software Development
7600 Boone Avenue North
Brooklyn Park, MN 55428
(W)612-424-4888
(WF)612-424-2853
nadya@network.com

Gary Ellis
Hewlett-Packard
Intelligent Networks Operation
1501 Page Mill Road
M/S 6U-O
Palo Alto, CA 94304
(W)415-857-8199
(WF)415-852-3762
garye@hpspd.spd.hp.com

Gus Elmer
AT&T Bell Laboratories
Room 1K-322
Crawfords Corner Road
Holmdel, NJ 07733
(W)908-949-8878
(WF)908-949-6689
`arch1!ae@att.com`

Fred Engel
Concord Communications, Inc.
Engineering
753 Forest Street
Marlboro, MA 01752
(W)508-460-4646
(WF)508-481-9772
`engel@concord.com`

Robert Enger
Contel Federal Systems
Software Engineering
15000 Conference Center Drive
Chantilly, VA 22021-3808
(W)703-818-5555
(WF)703-818-5485
`enger@seka.scc.com`

Johnny Eriksson
SUNET/NORDUnet
Operations Group
KTH
Stockholm, S-100 44
Sweden
(W)+46 8 790 6517
(WF)+46 8 24 11 79
`bygg@sunet.se`

Mike Erlinger
Micro Technology, Inc.
Network Products
5065 East Hunter Avenue
Anaheim, CA 92807
(W)714-970-0300
(W)800-999-9684
(WF)714-680-0279
`mike@mti.com`

Bill Fardy
Cabletron Systems
35 Industrial Way
Rochester, NH 03867
(W)603-332-9400
(WF)603-332-1019
`fardy@ctron.com`

Dino Farinacci
3Com Corporation
Internetworking Products
Building 300
5400 Bayfront Plaza
Santa Clara, CA 95052-8145
(W)408-764-5146
`dino@3com.com`

Elizabeth Feinler
SRI International
333 Ravenswood Drive
M/S EJ201
Menlo Park, CA 94025
(W)415-859-6287
(W)415-326-5512

Dennis Ferguson
University of Toronto
255 Huron Street
Toronto, ONTARIO M5S 1A1
Canada
(W)416-978-3328
(WF)416-978-6610
`dennis@canet.ca`

Charles Fineberg
Washington University
School of Medicine
4566 Scott Ave
M/S 8231
St. Louis, MO 63110
(W)(314) 362-3354
(WF)(314) 362-7183
fineberg@wums2.wustl.edu

Robert Fink
Lawrence Berkeley Laboratory
One Cyclotron Road
M/S 50B-2258
Berkeley, CA 94720
(W)415-486-5692
(WF)415-486-6363
rlfink@lbl.gov

Dale Finkelson
University of Nebraska
MIDnet
326 Administration
Lincoln, NE 68588-0496
(W)402-472-5108
(WF)402-472-5280
dmf@westie.unl.edu

Jeffrey Fitzgerald
FiberCom, Inc.
FDDI Development
P.O. Box 11966
Roanoke, VA 24022-1966
(W)703-342-6700
(WF)703-342-5961
jjf@fibercom.com

Peter Ford
Los Alamos National Laboratory
Advanced Computer Lab
M/S B287
Los Alamos, NM 87545
(W)505-665-0558
(WF)505-665-4939
peter@lanl.gov

Richard Fox
SynOptics Communications
950 Linden Avenue, #208
Sunnyvale, CA 94086
(W)408-764-1372
rfox@synoptics.com

Barbara Fraser
CERT
Software Engineering Institute
4500 Fifth Avenue
Pittsburgh, PA 15213-3890
(W)412-268-5010
(WF)412-268-6989
byf@cert.sei.cmu.edu

Ned Freed
Innosoft International, Inc.
250 W. First Street
Suite 240
Claremont, CA 91711
(W)714-624-7907
(WF)714-621-5319
net@ymir.claremont.edu

Daniel Friedman
Bolt, Beranek, and Newman
150 Cambridgepark Drive
M/S 20/5b
Cambridge, MA 02140
(W)617-873-3710
(WF)617-873-4986
danfriedman@bbn.com

Karen Frisa
Carnegie Mellon University
Network Development
4910 Forbes Avenue
M/S UCC 134
Pittsburgh, PA 15213
(W)412-268-8576
karen.frisa@andrew.cmu.edu

Tayang Fu
University of Maryland
5401 Wilkens Ave
Baltimore, MD 21228
(W)301-455-3964
(WF)301-455-3969
fu@algol.cs.umbc.edu

Vince Fuller
Stanford University
Networking & Communications Systems
Pine Hall 115
Stanford, CA 94305-4122
(W)415-723-6860
vaf@Standford.EDU

Charles Fumuso
Cray Research, Inc.
OSI Development
655F Lone Oak Drive
Eagan, MN 55121
(W)612-683-5449
(WF)612-683-5599
cwf@cray.com

Douglas Gale
National Science Foundation
NSFnet
1800 G Street, NW
Room 416
Washington, DC 20550
(W)202-357-9717
(WF)202-357-0320
dgale@note.nsf.gov

Shari Galitzer
The MITRE Corporation
7525 Colshire Drive
McLean, VA 22102-3481
(W)703-883-6229
(WF)703-883-7142
shari@gateway.mitre.org

Shawn Gallagher
Digital Equipment Corporation
153 Taylor Street
M/S TAY2-2/B4
Littleton, MA 01460
gallagher@quiver.enet.dec.com

Charles Gallucci
Netrix Corporation
13595 Dulles Technology Drive
Herndon, VA 22071
(W)703-742-6000
(WF)703-742-4048

James Galvin
Trusted Information Systems
3060 Washington Road
Route 97
Glenwood, MD 21738
(W)301-854-6889
(WF)301-854-5363
galvin@tis.com

Tony Genovese
Lawrence Livermore National
Laboratory
P.O. Box 5509
M/S L-561
Livermore, CA 94550
(W)415-423-2471
(WF)415-422-0435

Elise Gerich
University of Michigan
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)313-936-2120
(W)313-936-3000
(W)800-66-MERIT
(WF)313-747-3745
epg@merit.edu

Demi Getschko
Brazilian Research Network
R. Dr. Antonio Augusto De Almeida, 334
Campinas, S. Paulo 13083
Brazil
(W)+55 192 394141
(WF)+55 192 393070
"DEMI@FPSP.HEPNET"

Dave Geurs
Motorola, Inc.
Information Technology Management
1299 E. Algonquin Road
Schaumburg, IL 60196
(W)708-576-3970
(WF)708-576-6388
dgeurs@mot.com

Robert Gilligan
Sun Microsystems, Inc.
Internet Engineering
2550 Garcia Avenue
M/S MTV14-49
Mountain View, CA 94043-1100
(W)415-336-1012
(WF)415-336-3156
gilligan@sun.com

Joseph Golio
Cray Research, Inc.
655F Lone Oak Road
Eagan, MN 55121
(W)612-683-5536
(WF)612-683-5599
golio@cray.com

Fengmin Gong
Washington University
Dept. of Computer Science
One Brookings Drive
Campus Box 1045
St. Louis, MO 63130
(W)314-726-4163
lfg@wucs1.wustl.edu

Bret Gorsline
University of Michigan
721 E. Kingsley, #301
Ann Arbor, MI 48104
(W)313-995-4160
jinx@engin.umich.edu

Tom Grant
Xylogics, Inc.
53 Third Avenue
Burlington, MA 01803
(W)617-272-8140
(WF)617-273-5392
grant@xylogics.com

Fred Gray
BCSS/PSCNI
P.O. Box 9022
Building 4727
MSFC, AL 35824
(W)205-544-6101
(WF)205-544-4618
fred@homer.msfc.nasa.gov

Jeremy Greene
Coral Network Corporation
734 Forest Street
Marlboro, MA 01752
(W)508-460-6010
(WF)508-481-6258
greene@coral.com

Martin Gross
DCA/DCEC
Code DRFE
1860 Wiehle Avenue
Reston, VA 22090-5500
(W)703-437-2165
(WF)703-437-2104
gross@polaris.dca.mil

Phillip Gross
Advanced Network and Services
100 Clearbrook Road
Elmsford, NY 10523
(W)703-620-8990
(WF)703-620-0913
(W)914-789-5335
(W)914-789-5300
(WF)914-789-5310
pgross@nri.reston.va.us

Olafur Gudmundsson
University of Maryland
Department of Computer Science
College Park, MD 20742
(W)301-405-2767
(WF)301-405-6707
ogud@cs.umd.edu

Keith Hacke
Washington University
72 Frederick Lane
St. Louis, MO 63122
(W)314-362-4213
(WF)314-362-3454
hacke@informatics.wustl.edu

Robert Hagens
University of Wisconsin-Madison
Computer Science Department
1210 West Dayton Street
Madison, WI 53706
(W)608-262-5776
(WF)608-262-9777
(W)608-262-1017
hagens@cs.wisc.edu

Tony Hain
Lawrence Livermore National
Laboratory
P.O. Box 5509
M/S L-561
Livermore, CA 94550
(W)415-422-4200
(WF)415-422-0435
alh@eagle.es.net

Neil Haller
Bell Communications Research
445 South Street
M/S 2Q-280
Morristown, NJ 07962-1910
(W)201-829-4478
(WF)201-267-9298
nmh@bellcore.com

Martyne Hallgren
Cornell Theory Center
Networking Program
626 Eng. & Theory Center Bldg.
Ithaca, NY 14853-3801
(W)607-254-8767
(WF)607-254-8888
martyne@theory.tn.cornell.edu

Alf Hansen
University of Wisconsin-Madison
Computer Sciences Department
1210 West Dayton Street
Madison, WI 53706
(W)608-262-5084
(WF)608-262-9777
Alf.Hansen@pilot.cs.wisc.edu

Susan Hares
Merit Network, Inc.
Internet Engineering
1075 Beal Ave, NDSB
Ann Arbor, MI 48176
(W)313-936-2095
(WF)313-747-3745
skh@merit.edu

Eugene Hastings
Pittsburgh Supercomputer Center
4400 Fifth Avenue
Pittsburgh, PA 15213
(W)412-268-4960
(WF)412-268-5832
hastings@psc.edu

Frank Heath
Rockwell CMC
4311 Jamboree
Newport Beach, CA 92680
(W)714-833-6864
heath@cmc.com

Patrick Heisinger
McDonnell Douglas
Network Services
W126/3062182
P.O. Box 516
St. Louis, MO 63166
(W)314-232-6349
(WF)314-777-2006
heisinger@mdcguy.mdc.com

Sergio Heker
Princeton University
JVNVnet
B6von Neumann
Princeton, NJ 08544
(W)609-258-2411
(WF)609-258-2424
heker@jvnc.net

Robert Hinden
Bolt, Beranek, and Newman
Communications
10 Moulton Street
Cambridge, MA 02138
(W)617-873-3757
(WF)617-491-1496
hinden@bbn.com

Donald Hirsh
Meridian Technology Corporation
OEM Products
11 McBride Corporate Center Drive
Suite 250
Chesterfield, MO 63005
(W)314-532-7708
(WF)314-532-3242
hirsh@meridian.uucp

Russ Hobby
University of California
Computing Services
Surge II Room 1400
Davis, CA 95616
(W)916-752-0236
(WF)916-752-9154
rdhobby@ucdavis.edu

J. Paul Holbrook
CICNet, Inc.
2901 Hubbard, Pod G
Ann Arbor, MI 48105
(W)313-998-7680
(WF)313-998-6105
holbrook@cic.net

Jeffrey Honig
Cornell University
Information Technologies
143 Caldwell Hall
Ithaca, NY 14853-3801
(W)607-254-8768
(WF)607-254-8888
jch@devvax.tn.cornell.edu

Kathleen Huber
Bolt, Beranek, and Newman
Internet Switching and Security
Building 20, Suite 532
150 Cambridge Park Drive
Cambridge, MA 02140
(W)617-873-2520
(WF)617-873-4086
khuber@bbn.com

Steven Hunter
Lawrence Livermore National
Laboratory
Nat'l. Energy Res. Supercomputer Ctr.
P.O. Box 5509
M/S L-561
Livermore, CA 94551
(W)415-423-2219
(WF)415-422-0435
hunter@es.net

Joel Jacobs
The MITRE Corporation
Secure Distributed Systems
Burlington Road
M/S K302
Bedford, MA 01730
(W)617-271-7373
(WF)617-271-7722
jdj@mitre.org

Ole Jacobsen
Interop, Inc.
ConneXions
480 San Antonio Road
Suite 100
Mountain View, CA 94040
(W)800-INTEROP
(WF)415-949-1779
(W)415-941-3399
ole@csli.stanford.edu

B.V. Jagadeesh
3Com Corporation
5400 Bayfront Plaza
Santa Clara, CA 95052
(W)408-764-5169
(WF)408-764-5003
bvj@esd.3com.com

Mike Janson
Motorola, Inc.
Information Technology Management
1299 E. Algonquin Road
Schaumburg, IL 60196
(W)708-576-5403
(WF)708-576-6388
mjanson@mot.com

Dale Johnson
Merit Network, Inc.
Network Operations
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)313-936-2270
(W)313-936-3000
(WF)313-747-3745
dsj@merit.edu

William Johnston
Lawrence Berkeley Laboratory
M/S B50B-2239
Berkeley, CA 94720
(W)415-486-5014
(WF)415-486-6363
wejohnston@lbl.gov

Harold Jones
Digital Equipment Corporation
NPACE
550 King Street
M/S LKG2-2
Littleton, MA 01560
(W)508-486-5066
(WF)508-486-7417
hjones@nac.enet.dec.com

Kevin Jordan
Control Data Corporation
4201 Lexington Avenue, North
Arden Hills, MN 55126
(W)612-482-6835
(WF)612-482-2791
kej@udev.cdc.com

Dan Jordt
University of Washington
Networks and Distributed Computing
3737 Brooklyn Avenue, NE
Seattle, WA 98105
(W)206-543-7352
(WF)206-543-3909
danj@nwnet.net

Ajay Kachrani
Digital Equipment Corporation
4 Technology Drive
M/S DSG1-1/F8
Westford, MA 01886-4196
(W)(508) 635-8614
kachrani@regent.enet.dec.com

Mukesh Kacker
Sun Microsystems, Inc.
2550 Garcia Avenue
M/S 14-49
Mountain View, CA 94043
(W)415-336-4915
(WF)415-336-3156
mukesh@sun.com

Sanjay Kapoor
Washington University
6490 Enright, Apt. 402
St. Louis, MO 63130
kapoor@wuee1.wustl.edu

Philip Karn
Bell Communications Research
Data Communications Research
445 South Street
M/S 2Q-354
Morristown, NJ 07962-1910
(W)201-829-4299
(WF)201-267-9298
karn@thumper.bellcore.com

Frank Kastenholz
Clearpoint Research
Corporation
35 Parkwood Drive
Hopkinton, MA 01748
(W)508-235-2000
kasten@asherah.clearpoint.com

Neil Katin
Sun Microsystems, Inc.
Electronic Mail
2550 Garcia Avenue
Mt. View, CA 94043
(W)415-336-3718
(WF)415-962-9421
katin@eng.sun.com

Dave Katz
Merit Network, Inc.
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)313-936-3000
(WF)313-747-3745
dmk@merit.edu

Manu Kaycee
Digital Equipment Corporation
T&N/EIC Systems Engineering
Digital Drive
M/S MKO1-2/H10
Merrimack, NH 03054-9501
(W)603-884-5503
(WF)603-884-1036
kaycee@trlian.enet.dec.com

Douglas Kerr
Novell, Inc.
1340 Treat Boulevard
Walnut Creek, CA 94596
(W)415-975-4504
(WF)415-938-2562
dougk@wc.novell.com

Tom Kessler
Sun Microsystems, Inc.
Internet Engineering
2550 Garcia Avenue
M/S MTV14-49
Mountain View, CA 94043
(W)415-336-3145
(WF)415-336-3156
kessler@sun.com

Nam Keung
IBM
11400 Burnet Road
Austin, TX 78758
(W)512-823-4170

Kenneth Key
University of Tennessee
107 Agnes Hall
Knoxville, TN 37996
(W)615-974-5067
key@cs.utk.gdy

Charles Kimber
Dataproducts Corporation
6200 Canoga Avenue
Woodland Hills, CA 91365
(W)818-887-8009
(WF)818-887-8680

Jim Kinder
FiberCom, Inc.
FDDI Development
P.O. Box 11966
3353 Orange Avenue, NE
Roanoke, VA 24022-1966
(W)703-342-6700
(WF)703-342-5961
jdk@fibercom.com

Darren Kinley
RISQnet - CRIM
Suite 500
3744 Jean Brillant
Montreal, Quebec H3T 1P1
Canada
(W)514-340-5700
(WF)514-340-5777
kinley@crim.ca

Peter Kirstein
University College London
Department of Computer Science
Gower Street
London, WC1E 6BT
UK
(W)44-71-380-7286
(WF)44-71-387-1397
kirstein@cs.ucl.ac.uk

Anders Klemets
Swedish Institute of
Computer Science
CMU School of Computer Science
Pittsburgh, PA 15213
(W)412-268-7694
(WF)412-681-5739
klemets@cs.cmu.edu

Holly Knight
Apple Computer, Inc.
A/UX Development
10300 Bubb Road; Bubb 5
M/S 50-UX
Cupertino, CA 95015
(W)408-974-5110
(WF)408-974-0892
holly@apple.com

Jim Knowles
Sterling Software Federal Systems
NASA Ames Research Center
M/S 233-18
Moffett Field, CA 94035-1000
(W)415-604-3056
(WF)415-604-6999
jknowles@trident.arc.nasa.gov

Stev Knowles
FTP Software, Inc.
Engineering
26 Princess Street
Wakefield, MA 01880-3004
(W)617-246-0900
(WF)617-245-7943
stev@ftp.com

Shelly Knueven
Washington University
Office of the Network Coordinator
One Brookings Drive
Campus Box 1048
St. Louis, MO 63130-4899
(W)314-726-8524
(WF)314-726-7302
shelly@wugate.wustl.edu

Alex Koifman
Bolt, Beranek, and Newman
IP Routing
150 Cambridge Park Drive
M/S 031
Cambridge, MA 02140
(W)617-873-8610
(WF)617-873-4086
akoifman@bbn.com

Cheryl Krupczak
NCR Corporation
Network Products Division
NCRNet Manager Development
3245 Platt Springs Road
West Columbia, SC 29169
(W)803-739-7934
(WF)803-739-7745
clefor@secola.columbia.ncr.com

Tracy LaQuey Parker
University of Texas
Computation Center
M/S COM 1
Austin, TX 78712
(W)512-471-5046
(WF)512-471-1582
tracy@utexas.edu

Mary Louise Laier
Apple Computer, Inc.
20525 Mariani Avenue
M/S 35K
Cupertino, CA 95014
(W)408-974-6034
laierl@applelink.apple.com

Michelle Landriault
Bell Northern Research
P.O. Box 3511
Station C
Ottawa, Ontario K1Y 4H7
Canada
(W)613-763-5607
(WF)613-763-8581
crm57a@bnr.ca

Ruth Lang
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025
(W)415-859-5608
(WF)415-859-6028
rlang@nisc.sri.com

Nik Langrind
Shiva Corporation
Engineering
1 Cambridge Center
Cambridge, MA 02142
(W)617-252-6300
(WF)617-252-6852
nik@shiva.com

Vincent Lau
Sun Microsystems, Inc.
2550 Garcia Avenue
Mountain View, CA 94043
(W)415-336-1069
vlau@sun.com

Anthony Lauck
Digital Equipment Corporation
Telecommun. & Networks Architecture
550 King Street
M/S LKG1-2/E19
Littleton, MA 01460-1289
(W)508-486-7644
(W)508-486-7417
(WF)508-486-5279
lauck@tl.enet.dec.com

Walter Lazear
The MITRE Corporation
W-31 (Networking Center)
7525 Colshire Drive
M/S W420
McLean, VA 22102-3481
(W)703-883-6515
(WF)703-883-7142
lazear@gateway.mitre.org

Eliot Lear
Intelligenetics, Inc.
700 East El Camino Real
Mountain View, CA 94040
(W)415-962-7300
(W)415-962-7302
lear@net.bio.net

Dory Leifer
University of Michigan
Information Technology Division
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)313-763-4896
(WF)313-747-3745
Dory.Leifer@terminator.cc.umich.edu

Steven Lendt
University of Nebraska at Omaha
60th Dodge Street
Omaha, NE 68182-0051
(W)402-554-2523
(WF)402-554-3475
network@zeus.unomaha.edu

Mark Leon
NASA
Ames Research Center
Moffett Field
M/S 233 8
Mountain View, CA 94035
(W)415-604-6498
(WF)415-604-6999
leon@nsipo.arc.nasa.gov

Tony Li
cisco Systems
1525 O'Brien Drive
Menlo Park, CA 94025
(W)415-688-8245
(WF)415-688-8282
tli@cisco.com

Richard Libby
NCR Corporation
Network Products Division
Advanced Communications Adapters
2700 Snelling Avenue
M/S S010
St. Paul, MN 55113
(W)612-638-7255
(WF)612-638-7810
libby@c10sd6.stpaul.ncr.com

John Linn
Digital Equipment Corporation
Secure Systems Group
295 Foster Street
M/S LTN1-1/D07
Littleton, MA 01460-1123
(W)508-486-6588
(WF)508-486-6014
linn@zendia.enet.dec.com

David Lippke
University of Texas
Advanced Systems & Network Develop.
Academic Computing Center
P.O. Box 830688
Richardson, TX 75083-0688
(W)214-690-2632
(WF)214-690-2966
lippke@utdallas.edu

Mike Little
Bell Communications Research
RRC1N2119
444 Hoes Lane
Piscataway, NJ 08855
(W)908-699-2819
(WF)908-463-2078
little@ctt.bellcore.com

Joshua Littlefield
Cayman Systems, Inc.
26 Landsdowne Street
Cambridge, MA 02139
(W)617-494-1999
(WF)617-494-9270
josh@cayman.com

Then Liu
3Com Corporation
5400 Bayfront Plaza
Santa Clara, CA 95052
(W)408-764-5391
(WF)408-764-5089

John LoVerso
Concurrent Computer Corporation
One Technology Way
Westford, MA 01886
(W)508-392-2440
(WF)508-392-2492
loverso@westford.ccur.com

Daniel Long
Bolt, Beranek, and Newman
NEARnet Project
10 Moulton Street
M/S 6/3B
Cambridge, MA 02174
(W)617-873-2766
(WF)617-873-3776
long@nic.near.net

E. Paul Love
San Diego Supercomputer Center
Peripherals & Communications
P.O. Box 85608
San Diego, CA 92186
(W)619-534-5043
(WF)619-534-5152
loveep@sdsc.edu

Ron Mackey
Distributed Systems
International, Inc.
610 S. Knollwood Drive
Wheaton, IL 60187
(W)708-665-4693

Carl Malamud
Consultant
3230 Folsom Avenue
Boulder, CO 80304
(W)303-440-7052
(WF)303-440-7526
carl@malamud.com

Andrew Malis
Bolt, Beranek, and Newman
Switching Product Development
150 CambridgePark Drive
M/S 20/7a
Cambridge, MA 01240
(W)617-873-3419
(WF)617-873-4086
malis@bbn.com

Gary Malkin
FTP Software, Inc.
26 Princess Street
Wakefield, MA 01880-3004
(W)617-246-0900
(WF)617-245-7943
gmalkin@ftp.com

Allison Mankin
The MITRE Corporation
7525 Colshire Drive
M/S W425
McLean, VA 22102-3481
(W)703-883-7907
(WF)703-883-7142
mankin@gateway.mitre.org

Bill Manning
Texas Instruments, Inc.
6550 Chase Oaks
M/S 693
Plano, TX 74023
(W)713-274-2441
(WF)713-274-3911
bmanning@houston.sc.ti.com

Mike Marcinkevicz
California State University
Computer Center
800 N. State College Boulevard
Fullerton, CA 92634
(W)714-773-3009
(WF)714-773-3990
mdm@calstate.edu

Keith McCloghrie
Hughes LAN Systems
Engineering
1225 Charleston Road
Mountain View, CA 94043
(W)415-966-7934
(WF)415-960-3738
kzm@hls.com

Glenn McGregor
Merit Network, Inc.
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)313-936-3000
(WF)313-747-3745
ghm@merit.edu

Leo McLaughlin
FTP Software, Inc.
26 Princess Street
Wakefield, MA 01880
(W)617-246-0900
(WF)617-245-7943
ljm@ftp.com

Donna McMaster
SynOptics Communications, Inc.
4401 Great America Parkway
P.O. Box 58185
Santa Clara, CA 95051-8185
(W)408-764-1206
(WF)408-764-1299
dmcmaster@synoptics.com

Milo Medin
Sterling Software Federal Systems
Systems & Scientific Division
Federal Systems Group
1121 San Antonio Road
Palo Alto, CA 94303
(W)415-604-6440
(W)415-964-9900
(WF)415-969-3821
medin@nsipo.nasa.gov

Robert Meierhofer
NCR Corporation
Network Product Division
2700 N. Snelling Avenue
M/S S010
St. Paul, MN 55113
(W)612-638-7278
(WF)612-638-7810
bobm@cl0sd6.stpaul.ncr.com

Carol Melowitz
McDonnell Douglas
325 McDonnel Boulevard
Hazelwood, MD 63042
(W)301-232-5195
(WF)301-777-2006
melowitz@mdcgwy.mdc.com

Jay Melvin
Comtech Labs, Inc.
4005 Miranda Avenue, #150
Palo Alto, CA 94304-1218
(W)415-424-4500
(WF)415-424-1214
infopath@well.sf.ca.us

Linda Melvin
Comtech Labs, Inc.
4005 Miranda Avenue, #150
Palo Alto, CA 94304-1218
(W)415-813-4500
(WF)415-424-1214
infopath@well.sf.ca.us

Alexander Merola
Lawrence Berkeley Laboratory
Information and Computing
One Cyclotron Road
Berkeley, CA 94720
(W)415-486-7440
(W)415-486-6363
axmerola@lbl.gov

April Merrill
Department of Defense
9800 Savage Road
Ft. Meade, MD 10755-6000
(W)301-684-7302

Donald Merritt
Ballistics Research Laboratory
Sys. Engr./Concepts Analysis Div.
Network and Computer Systems Team
SLCBR-Se-A
Aberdeen Proving, MD 21005-5066
(W)301-278-6808
(WF)301-278-5075
don@brl.mil

David Miller
The MITRE Corporation
Burlington Road
D115
M/S E066
Bedford, MA 01730
(W)617-271-3993
dtm@ulana.mitre.org

Cyndi Mills
Bolt, Beranek, and Newman
Network Analysis
150 Cambridge Park Drive
M/S 20/608
Cambridge, MA 02140
(W)617-873-4143
(WF)617-873-8202
cmills@bbn.com

Greg Minshall
Novell, Inc.
Suite 500
1340 Treat Boulevard
Walnut Creek, CA 94596
(W)415-975-4507
(WF)415-947-1238
minshall@wc.novell.com

Lynn Monsanto
Sun Microsystems, Inc.
2550 Garcia Avenue
M/S 14-49
Mountain View, CA 94043
(W)415-336-2946
(WF)415-336-3156
monsanto@eng.sun.com

Mark Moody
MOREnet
200 Hienkel Building
Columbia, MO 65211
(W)314-882-2000
(WF)314-882-9975
ccmarkm@umcvmb.missouri.edu

Keith Moore
University of Tennessee
107 Ayres Hall
Knoxville, TN 37996
(W)615-546-4566
moore@cs.utk.edu

Robert Morgan
Stanford University
115 Pine Hall
M/S 4122
Stanford, CA 94305-4122
(W)415-723-9711
(WF)415-723-0010
morgan@jessica.stanford.edu

Dennis Morris
Defense Communications
Engineering Center
1860 Wiehle Avenue
Reston, VA 22090-5500
(W)703-437-2584
(WF)703-437-2104
morrisd@imo-uvax.dca.mil

Donald Morris
National Center for
Atmospheric Research
Scientific Computing Division
P.O. Box 3000
Boulder, CO 80307-3000
(W)303-497-1282
(WF)303-497-1137
morris@ucar.edu

John Morrison
Los Alamos National Laboratory
Computer Division
Group C-5
M/S B255
Los Alamos, NM 87545
(W)505-667-3310
(WF)505-665-4361
jfm@lanl.gov

Candice Moshos
Corporation for National
Research Initiatives
1895 Preston White Drive
Suite 100
Reston, VA 22091
(W)703-620-8990
cmoshos@nri.reston.va.us

James Mostek
Cray Research, Inc.
655F Lone Oak Drive
Eagan, MN 55121
(W)612-683-5470
(WF)612-683-5599
mostek@cray.com
(W)612-683-5470
(WF)612-683-5599
mostek@cray.com

John Moy
Proteon, Inc.
Two Technology Drive
Westborough, MA 01581-1799
(W)508-898-2800
(WF)508-366-8901
jmoy@proteon.com

Gary Mussar
Bell Northern Research
P.O. Box 3511
Station C
Ottawa, Ontario K1Y 4H7
CANADA
(W)613-763-4937
(WF)613-763-8581
mussar@bnr.ca

Chris Myers
Washington University
Office of the Network Coordinator
Campus Box 1048
One Brookings Drive
St. Louis, MO 63130-4899
(W)314-726-7390
(WF)314-726-7302
chris@wugate.wustl.edu

Raj Nair
University of Maryland
Department of Computer Science
5401 Wilkens Avenue
Baltimore, MD 21228
(W)301-455-3966
(WF)301-455-3969
nair@algol.cs.umbc.edu

Mark Needleman
University of California
Division of Library Automation
Kaiser Center, Room 854
300 Lakeside Drive
Oakland, CA 94612-3550
(W)415-987-0530
(WF)415-839-3573
mhn@stubbs.ucop.edu

Andy Nicholson
Cray Research, Inc.
655F Lone Oak Drive
Eagan, MN 55121
(W)612-683-5473
(WF)612-683-5599
droid@cray.com

Bill Norton
Merit Network, Inc.
1075 Beal Ave.
Ann Arbor, MI 48109-2112
(W)313-936-2656
wbn@merit.edu

David O'Leary
SURAnet
8400 Baltimore Boulevard
College Park, MD 20740
(W)301-982-3214
(WF)301-982-4605
oleary@sura.net

Alan Oppenheimer
Apple Computer, Inc.
20525 Mariani Avenue
Cupertino, CA 95014
(W)408-974-4744
(W)408-996-1010
oppenheimer1@applelink.apple.com

Christos Papadopoulous
Washington University
Dept. of Computer Science
Box 1045
One Brookings Drive
St. Louis, MO 63130
(W)314-889-5087
(WF)314-645-5874
christos@dworkin.wustl.edu

Philippe Park
Bolt, Beranek, and Newman
Systems and Technologies
10 Moulton Street
Cambridge, MA 02138
(W)617-873-2892
ppark@bbn.com

Brad Parker
Cayman Systems, Inc.
26 Landsdowne Street
Cambridge, MA 02139
(W)617-494-1999
(WF)617-494-9270
brad@cayman.com

Gurudatta Parulkar
Washington University
Computer Science Department
Bryan Hall
1 Brookings Drive, Box 1045
St. Louis, MO 63130
(W)314-889-4621
(WF)314-726-7302
guru@flora.wustl.edu

Michael Patton
Massachusetts Institute of
Technology
Laboratory for Computer Science
545 Technology Square
Cambridge, MA 02139
(W)617-253-6061
(WF)617-258-8682
map@lcs.mit.edu

Robert Peglar
Network Systems, Inc.
Internetwork Group
7600 Boone Avenue North
Minneapolis, MN 55428
(W)612-424-4888
robp@anubis.network.com

David Perkins
3Com Corporation
5400 Bayfront Plaza
P.O. Box 58145
Santa Clara, CA 95052-8145
(W)408-764-5983
(WF)408-764-6135
dave_perkins@3com.com

Richard Pethia
Carnegie Mellon University
Software Engineering Institute
4500 Fifth Avenue
Pittsburgh, PA 15213-3890
(W)412-268-7739
(WF)412-268-6989
rdp@cert.sei.cmu.edu

James Philippou
Xyplex, Inc.
330 Codmanhill Road
Boxborough, MA 01719
(W)508-264-9900
(WF)508-264-9930
japhilippou@eng.xyplex.com

Mark Pohlman
Washington University
Office of the Network Coordinator
Campus Box 1048
St. Louis, MO 63130-4899
(W)314-726-7391
(WF)314-726-7302
mark@wugate.wustl.edu

Ron Poppen-Chambers
Hewlett-Packard: CND
3404 East Harmony Road
Fort Collins, CO 80525
(W)303-229-2842
rpc@hpcnd.cnd.hp.com

Rehmi Post
FTP Software
26 Princess Street
Wakefield, MA 01880
(W)617-863-9544
rehmi@ftp.com

Stephanie Price
Rockwell - CMC
125 Cremona
Santa Barbara, CA 93105
(W)805-968-4262
(WF)805-968-8256
price@cmc.com

K.K. Ramakrishnan
Digital Equipment Corporation
550 King Street
M/S LKG1-2/A19
Littleton, MA 01754
(W)508-486-7267
rama@kalvi.enet.dec.com

James Reeves
SynOptics Communications
FDDI Software Development
P.O. Box 58185
4401 Great American Parkway
Santa Clara, CA 95052-8185
(W)408-764-1440
(WF)408-988-5525
jreeves@synoptics.com

Michael Reilly
Digital Equipment Corporation
Network Systems Lab
505 Hamilton Avenue
Suite 101
Palo Alto, CA 94301
(W)415-688-1323
(W)415-688-1320
(WF)415-324-2797
reilly@pa.dec.com

John Reinart
Cray Research, Inc.
655F Lone Oak Drive
Eagan, MN 55121
(W)612-683-5329
(WF)612-683-5599
reinart@cray.com
(W)612-683-5329
(WF)612-683-5599
reinart@cray.com

Yakov Rekhter
IBM Corporation
T.J. Watson Research Center
P.O. Box 218, Route 134
Yorktown Heights, NY 10598
(W)914-945-3896
yakov@ibm.com

Joel Replogle
National Center for
Supercomputing Applications
605 East Springfield Avenue
Champaign, IL 61820
(W)217-244-0636
(WF)217-244-1987
replogle@ncsa.uiuc.edu

Robert Reschly
Ballistics Research Laboratory
ATTN: SLCBR-SE-A (RESCHLY)
Aberdeen Proving, MD 21005-5066
(W)301-278-6808
(WF)301-278-5075
`reschly@brl.mil`

Joyce Reynolds
Information Science Institute
University of Southern California
4676 Admiralty Way, Suite 1001
Marina del Rey, CA 90292-6695
(W)213-822-1511
(WF)213-823-6714
`jkrey@isi.edu`

Anil Rijsinghani
Digital Equipment Corporation
153 Taylor Street
M/S TAY2-2/B4
Littleton, MA 01460
(W)508-952-3520
`anil@levers.enet.dec.com`

Ron Roberts
Stanford University
115 Pine hall
Room 175-B
Stanford, CA 94305-4122
(W)415-723-3352
(W)415-723-3909
(WF)415-723-0010
`roberts@jessica.stanford.edu`

Kary Robertson
Concord Communications, Inc.
753 Forest Street
Marlborogh, MA 01752
(W)508-460-4646
(WF)508-481-9772
`kr@concord.com.kr`

Manoel Rodrigues
AT&T Bell Laboratories
Data Network Analysis
Room 1F-412
Crawfords Corner Road
Holmdel, NJ 07733
(W)908-949-4655
`manoel.rodrigues@att.com`

Allyn Romanow
Sun Microsystems, Inc.
2550 Garcia Avenue
Mountain View, CA 94043
(W)415-336-1326
`allyn@eng.sun.com`

Karen Roubicek
Bolt, Beranek, and Newman
Distributed Systems
10 Moulton Street
M/S 6/3B
Cambridge, MA 02138
(W)617-873-3361
(WF)617-873-3776
`roubicek@bbn.com`

Allan Rubens
Merit Network, Inc.
1075 Beal Ave
Ann Arbor, MI 48109
(W)313-747-0417
(WF)313-747-3745
`acr@merit.edu`

Gregory Ruth
Bolt, Beranek, and Newman
150 Cambridge Park Drive
Cambridge, MA 02140
(W)617-873-3150
(WF)617-873-8202
`gruth@bbn.com`

Jan Michael Rynning
NADA
KTH
Stockholm, S-100 44
Sweden
(W)+46 8 790 6288
(WF)+46 8 790 0930
jmr@nada.kth.se

Ray Samora
Proteon, Inc.
Engineering
Two Technology Drive
Westborough, MA 01581
(W)508-898-2800
(WF)508-366-8901
rvs@proteon.com

George Sanderson
McDonnell Douglas
P.O. Box 516
M/S 080 1480
St. Louis, MO 63166
(W)314-895-5342
(WF)314-895-5297
sanderson@mdc.com

Jonathan Saperia
Digital Equipment Corporation
40 Old Bolton Road
M/S OGO1-1/J17
Stow, MA 01775
(W)508-496-8333
saperia@decwrl.enet.dec.com

Greg Satz
cisco Systems
Network Products
1525 O'Brien Drive
Menlo Park, CA 94025
(W)415-688-8245
(WF)415-688-8282
satz@cisco.com

Mark Schaefer
DAVID SYSTEMS
701 E. Evelyn Avenue
Box 3718
Sunnyvale, CA 94088-3718
(W)408-720-6804
(WF)408-720-1337
schaefer@davidsys.com

Jeffrey Schiller
Massachusetts Institute of
Technology
Network Services
Room E40-311
1 Amherst Street
Cambridge, MA 02139
(W)617-253-8400
(W)617-253-0161
(WF)617-258-8736
jis@mit.edu

Tim Seaver
Microelectronics Center
of North Carolina
Center for Microelectronics
P.O. Box 12889
3021 Cornwallis Road
Research Triangle Park, NC 27709-2889
(W)919-248-1973
(WF)919-248-1455
tas@mcnc.org

Paul Selkirk
FTP Software, Inc.
26 Princess Street
Wakefield, MA 01880
(W)617-426-0900
(WF)617-245-7943
paul@ftp.com

Marc Sheldon
University of Dortmund
IRB
August-Schmidt-Strasse 12
P.O. Box 500500
D-W-4600, Dortmund 50
Germany
(W)49 231 7552444
(WF)49 231 7552386
ms@Germany.eu.net

Daisy Shen
IBM Corporation
Advance Communications
T.J. Watson Research Center
P.O. Box 218
Yorktown Heights, NJ 10598
(W)914-945-3632
(WF)914-945-4437
daisy@watson.ibm.com

Jim Sheridan
IBM Corporation
NSFnet
PO Box 334
Whitmore Lake, MI 48189
(W)313-393-6537
(WF)313-747-3745
jsherida@ibm.com

Mark Sherman
Information Technology Center
Carnegie Mellon University
4910 Forbes Avenue
Pittsburgh, PA 15213
(W)412-268-6727
(WF)412-268-6787
mss+@andrew.cmu.edu

Steven Sherry
Xyplex, Inc.
330 Codman Hill Road
Boxborough, MA 01719-1708
(W)508-264-9900
(WF)508-264-9930
shsherry@eng.xyplex.com

Stephen Shew
Bell Northern Research
P.O. Box 3511
Station C
Ottawa, Ontario K1Y 4H7
CANADA
(W)613-763-2462
(WF)613-763-8581
sdshew@bnr.ca

Dror Shindelman
Fiberonics International
Advanced Technology Center
Haifa, 31905
ISRAEL
(W)972-4-313313
(WF)972-4-536360
pbrenner@sparta.spartacus.com

Deepinder Sidhu
University of Maryland
Department of Computer Science
5401 Wilkens Avenue
Baltimore, MD 21228
(W)301-455-3028
(WF)301-455-3969
sidhu@umbc3.umbc.edu

Anil Singhal
Frontier Software Development, Inc.
1501 Main Street, #40
Tewksbury, MA 01876
(W)508-851-8872
(WF)508-851-6956

Ursula Sinkewicz
Digital Equipment Corporation
550 King Street
M/S LKG1-3/L06
Littleton, MA 01460-1289
(W)508-486-5604
(WF)508-486-7417
sinkewic@decvax.dec.com

Dana Sitzler
Merit Network, Inc.
Network Development & Support Bldg.
1075 Beal Avenue
Ann Arbor, MI 48109-2112
(W)313-936-3000
(WF)313-747-3745
dds@merit.edu

Sam Sjogren
TGV
Engineering
603 Mission Street
Santa Cruz, CA 95060
(W)408-427-4366
(WF)408-427-4365
sjogren@tgv.com

Patricia Smith
Merit Network, Inc.
1075 Beal Avenue - NDSB
Ann Arbor, MI 48109-2112
(W)313-936-3000
(WF)313-747-3745
psmith@merit.edu

Lawrence Snodgrass
CREN
1112 16th Street, NW
Suite 600
Washington, DC 20036
(W)202-872-4200
(WF)202-872-4318
snodgras@educom.edu

Frank Solensky
Clearpoint Research
Corporation
35 Parkwood Drive
Hopkinton, MA 01748
(W)508-235-2000
solensky@clearpoint.com

Evan Solley
Farallon Computing, Inc.
4730 SW Macadam Avenue
Portland, OR 97201
(W)503-292-1374
(WF)503-226-6543
solley@applelink.apple.com

Brad Solomon
Boeing Computer Support Service
P.O. Box 9022
MSFC
MSFC, AL 35812
(W)205-544-6762
(WF)205-544-4618
bsolomon@hobbes.msfc.nasa.gov

Albert Soule
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213
(W)412-268-6821
(WF)412-268-5758
als@sei.cmu.edu

Mike Spengler
Minnesota Supercomputer Center
1200 Washington Avenue South
Minneapolis, MN 55415
(W)612-626-1757
mks@msc.edu

Michael St. Johns
Department of Defense
Attn: T71
9800 Savage Road
Ft. George Meade, MD 20755
(W)301-688-6741
stjohns@umd5.umd.edu

Michael Stanton
Brazilian Research Network
R. Dr. Antonio Augusto de Almeida, 334
Campinas, S. Paulo 13083
Brazil
(W)+55 192 394141
(WF)+55 192 393070
"usermas@lncc.bitnet"

Martha Steenstrup
Bolt, Beranek, and Newman
Network Analysis
Room 20/665
150 Cambridge Park Drive
Cambridge, MA 02140
(W)617-873-3192
(WF)617-873-8202
msteenst@bbn.com

James Sterbenz
Washington University
Bryan 405, Box 1115
St. Louis, MO 63130-4636
(W)314-726-4203
(WF)314-726-7302
jps@dworkin.wustl.edu

Bob Stewart
Xyplex, Inc.
Architecture & Technology
330 Codman Hill Road
Boxborough, MA 01719
(W)508-264-9900
(WF)508-264-9930
rlstewart@eng.xyplex.com

Bernhard Stockman
SUNET/NORDUnet
KTH
Stockholm, S-100 44
Sweden
(W)46 8 790 6519
(WF)46 8 241 179
boss@sunet.se

Roxanne Streeter
Sterling Software Federal Systems
Science Internet Project Office
M/S 233-8
Moffett Field, CA 94035
(W)415-604-4845
(WF)415-604-0063
streeter@nsipo.nasa.gov

Allen Sturtevant
Lawrence Livermore National
Laboratory
PO Box 5509
M/S L-561
Livermore, CA 94550
(W)415-422-8266
(WF)415-422-0435
sturtevant@ccc.nmfecc.gov

Mitch Sukalski
Los Alamos National Laboratory
Group C-4
M/S B254
Los Alamos, NM 87545
(W)505-665-1637
(WF)505-665-4361
mws@lanl.gov

Natee Tanchanpongs
Washington University
Office of the Network Coordinator
One Brookings Drive
Campus Box 1048
St. Louis, MO 63130-4899
(W)314-726-7392
(WF)314-726-7302
natu@wugate.wustl.edu

Marty Tanner
Southwestern Bell Telephone
One Bell Center
St. Louis, MO 63101

Kaj Tesink
Bell Communications Research
NVC 1H318
331 Newman Springs Road
Red Bank, NJ 07748
(W)908-758-5254
(WF)908-530-6862
kaj@nvuxr.cc.bellcore.com

Michael Tharenos
Stanford University
Show Net Operations
480 San Antonio
Mountain View, CA 94305
(W)415-917-2221
tharenos@jessica.stanford.edu

Joanie Thompson
Sterling Software Federal Systems
NASA Sciences Internet
NASA/Ames Research Center
M/S 233-8
Moffett Field, CA 94035-1000
(W)415-604-4550
(WF)415-604-0063
joanie@nsipo.nasa.gov

Dean Throop
Data General Corporation
Data Communications
62 Alexander Drive
Research Triangle, NC 27709
(W)919-549-8421
(WF)919-248-5942
throop@dg-rtp.dg.com

Claudio Topolcic
Corporation for National
Research Initiatives
Internet Development
1895 Preston White Drive
Suite 100
Reston, VA 22091
(W)703-620-8990
(WF)703-620-0913
topolcic@nri.reston.va.us

William Townsend
Xylogics, Inc.
53 Third Avenue
Burlington, MA 01803
(W)617-272-8140
(WF)617-273-5392
townsend@xylogics.com

Glenn Trewitt
Digital Equipment Corporation
Network Systems Laboratory
505 Hamilton Avenue
Suite 101
M/S UCH-1
Palo Alto, CA 94301
(W)415-688-1324
(WF)415-324-2797
trewitt@pa.dec.com

Paul Tsuchiya
Bell Communications Research
435 South Street
M/S MRE 2P-281
Morristown, NJ 07962-1910
(W)201-829-4484
(W)201-267-9298
tsuchiya@bellcore.com

Maurice Turcotte
Racal-Milgo Information Systems
Suite 300
400 Embassy Row
Atlanta, GA 30328
(W)404-804-5190
(WF)404-804-5250
dnmrt@rm1.uucp

Mike Turico
Motorola - Information
Technology Managment
1299 E. Algonquin Road
Schaumburg, IL 60196
(W)708-578-8409
(WF)708-576-6388
mturico@mot.com

Jonathan Turner
Washington University
Computer and Comm. Research Center
One Brookings Drive, Campus Box 1115
St. Louis, MO 63130
(W)314-889-6193
(W)314-726-4434
jst@wucs1.wustl.edu

Robert Ullmann
Prime Computer, Inc.
500 Old Connecticut Path
M/S 10-30
Framingham, MA 01701
(W)508-620-2800
(WF)508-620-2800
ariel@relay.prime.com

Gregory Vaudreuil
Corporation for National
Research Initiatives
Internet Engineering
1895 Preston White Drive, Suite 100
Reston, VA 22091
(WF)703-620-0913
(W)703-620-8990
gvaudre@nri.reston.va.us

Ross Veach
University of Illinois
Computing Services Office
1304 West Springfield Avenue
Urbana, IL 61801
(W)217-244-4274
(WF)217-244-7089
rrv@uiuc.edu

John Veizades
Apple Computer, Inc.
Networking and Communications
20525 Mariani Avenue
M/S 35-K
Cupertino, CA 95014
(W)408-974-2672
(WF)408-974-9083
veizades@apple.com

Sudhanshu Verma
Hewlett-Packard
Intelligent Networks Operation
1501 Page Mill Road
M/S 6U-O
Palo Alto, CA 94304
(W)418-857-8196
(WF)415-852-3762
verma@hpindbu.cup.hp.com

Steven Waldbusser
Carnegie Mellon University
Computing and Communications
4910 Forbes Avenue
UCC Building
Pittsburgh, PA 15213-3890
(W)412-268-6628
(WF)412-268-5249
steven.waldbusser@andrew.cmu.edu

Francis Waldman
AT&T Bell Laboratories
Crawfords Corner Road
M/S 1F337A
Holmdel, NJ 07733
(W)908-949-3096
(WF)908-949-5775
feri@arch2.att.com

Justin Walker
Apple Computer, Inc.
10300 Bubb Road
M/S 50-UX
Cupertino, CA 95014
(W)408-974-4427
(WF)408-974-0892
justin@apple.com

Shujiuan Wang
IBM
472 Wheelers Farms Road
Milford, CT 06410
(W)203-250-9459
swang@ibm.com

Carol Ward
University of Colorado
Westnet
3645 Marine Street
M/S 0455
Boulder, CO 80027
(W)303-492-5860
(WF)303-492-4198
cward@spot.colorado.edu

Chris Weider
Merit Network, Inc.
1075 Beal Ave
Ann Arbor, MI 48109-2112
(W)313-936-3000
(W)313-936-2090
(WF)312-747-3745
clw@merit.edu

Daniel Weidman
Washington University
Department of Surgery
Campus Box 8109
4480 Clayton Avenue
St. Louis, MO 63110
(W)314-362-7683
(WF)314-531-7325
weidman@wudos2.wustl.edu

Jonathan Wenocur
Shiva Corporation
Engineering
1 Cambridge Center
Cambridge, MA 02142
(W)617-252-6300
(WF)617-252-6852
jhw@shiva.com

Kathleen Wilde
Digital Equipment Corporation
Telecommunications & Networking
110 Spitbrook Road
Nashua, NH 03062-2698
(W)603-881-0387
(WF)508-486-5279
wilde@decvax.dec.com

Rick Wilder
The MITRE Corporation
7525 Colshire Drive
M/S W425
McLean, VA 22102-3481
(W)703-883-6174
(WF)703-883-7142
rick@gateway.mitre.org

Bert Williams
Synernetics
Western Operations
800 West El Camino Real
Suite 180
Mountain View, CA 94040
(W)415-903-2225
(WF)415-903-2237
bert@mailgate.synnet.com

Walter Wimer
Carnegie Mellon University
Network Development
4910 Forbes Avenue
UCC Building
Pittsburgh, PA 15213-3890
(W)412-268-6252
(WF)412-268-5249
walter.wimer@andrew.cmu.edu

Linda Winkler
Argonne National Laboratory
Computing and Telecommunications
Building 221, B-251
9700 South Cass Avenue
M/S 221/B251
Argonne, IL 60439
(W)708-910-6104
(W)708-972-7236
(WF)708-972-5983
lwinkler@anl.gov

Cathy Wittbrodt
Lawrence Livermore National
Laboratory
Nat'l Engergy Research Supercomp.
P.O. Box 5509
M/S L-561
Livermore, CA 94550
(W)415-466-4016
(WF)415-422-0435
cjw@nersc.gov

Steve Witten
Hewlett-Packard
1501 Page Mill Road
M/S 6U-0
Palo Alto, CA 94304
(W)415-857-5569
(WF)415-852-3762

John Wobus
Syracuse University
Computing and Network Services
219 Machinery Hall
Syracuse, NY 13244
(W)315-443-4324
(WF)315-443-1819
jmwobus@suvm.acs.syr.edu

Stephen Wolff
National Science Foundation
Network & Communication Research
1800 G Street NW, Room 533
Washington, DC 20550
(W)202-357-9717
(W)202-357-0320
steve@nsf.gov

Wing Fai Wong
Salomon Technology Services, Inc.
745 Route 3
Rutherford, NJ 07070
(W)201-896-7322
(WF)201-460-0674
wfwong@malta.sbi.com

C. Philip Wood
Los Alamos National Laboratory
Network Engineering
Group C-5
M/S B255
Los Alamos, NM 87545
(W)505-667-2598
(WF)505-665-4361
cpw@lanl.gov

Research
and
Development
for
the
National
Information
Infrastructure